IBM FAStT Storage Manager Version 8.21 for
UNIX and AIX Environments
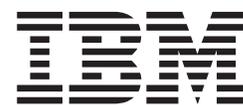
IBM

Installation and Support Guide

IBM FAStT Storage Manager Version 8.21 for
UNIX and AIX Environments

IBM

Installation and Support Guide

> **Note:**
> Before using this information and the product it supports, read the information in "Notices" on page 77.

Contents

Figures

Tables

About this document

This document provides information about setting up, installing, configuring, and working with the IBM Fibre Array Storage Technology (FAStT) Storage Manager Version 8.21 in UNIX® and AIX™ environments.

Throughout this document, the terms *storage management software* and *Storage Manager 8.21* refer to the IBM FAStT Storage Manager Version 8.21. Individual components of the storage management software are identified by name.

Use this document to:
- Determine the hardware and software that is required to install the management software into your subsystem network
- Integrate the necessary hardware components
- Install the management software
- Upgrade controller NVSRAM and firmware
- Identify management features that are unique to your specific installation

Who should read this document

This document is intended for system administrators and storage administrators who are responsible for installing software. Readers should have knowledge of RAID, SCSI, and fibre-channel technology, and should also have working knowledge of the applicable operating systems that are used with the management software.

How this document is organized

Chapter 1, "Introduction" on page 1 provides an introduction to the Storage Manager 8.21 product, including information about product resources in addition to this document.

Chapter 2, "Installing storage management station software on AIX systems" on page 11 provides information about installing and using Storage Manager 8.21 with an AIX operating system, step-by-step instructions for installing the AIX software on a management station, and step-by-step instructions for installing the AIX management software on a host.

Chapter 3, "Installing storage management station software on HP-UX systems" on page 29 provides information about installing and using Storage Manager 8.21 with the HP-UX operating system, step-by-step instructions for installing the HP-UX software on a management station, and step-by-step instructions for installing the HP-UX management software on a host.

Chapter 4, "Installing storage management station software on Solaris systems" on page 43 provides information about installing and using Storage Manager 8.21 with the Solaris operating system, step-by-step instructions for installing the Solaris software on a management station, and step-by-step instructions for installing the Solaris management software on a host.

Chapter 5, "Completing the software installation" on page 55 provides step-by-step instructions for performing post-installation tasks.

Chapter 6, "Setting up the storage subsystem in a heterogeneous environment" on page 65 provides step-by-step instructions for setting up and running Storage Manager 8.21 in a heterogeneous environment.

Chapter 7, "Storage Manager 8.21 with high-availability cluster services" on page 67 provides information about the high-availability clustering system.

Appendix A, "AIX system requirements" on page 69, Appendix B, "HP-UX system requirements" on page 71, and Appendix C, "Solaris system requirements" on page 73 list the minimum hardware and software requirements that AIX, HP-UX, and Solaris systems, respectively, must meet to be used in a storage system with Storage Manager 8.21.

Appendix D, "MC Service Guard configuration details" on page 75 explains the necessity and procedure for correcting the primary and alternate paths of the imported volume groups changed after using vgimport -m -s with LVM commands.

Related publications

This section lists documents, online help systems, and Web sites. These resources provide related information that might be of interest to you.

Documents

The following documents are available in Adobe Acrobat PDF on the Storage Manager 8.21 installation CD and at the following Web site:

www.ibm.com/pc/support/

- *IBM FAStT Storage Manager Version 8.21 for UNIX and AIX Environments, Installation and Support Guide* (this document)
- *IBM FAStT Storage Manager Version 8.21 Installation and Support Guide for Windows NT® and Windows® 2000*
- *IBM FAStT Storage Manager Version 8.21 Installation and Support Guide for Novell NetWare*
- *IBM FAStT Storage Manager Version 8.21 Installation and Support Guide for Linux*
- *IBM TotalStorage™ FAStT700 Fibre Channel Storage Server Installation Guide*
- *IBM TotalStorage FAStT700 Fibre Channel Storage Server User's Guide*
- *IBM FAStT700 Fibre Channel Cabling Instructions*
- *IBM FAStT500 RAID Controller Enclosure Unit Installation Guide*
- *IBM FAStT500 RAID Controller Enclosure Unit User's Reference*
- *IBM Netfinity® Fibre Channel Cabling Instruction*
- *IBM FAStT200 and FAStT200 HA Storage Servers Installation and User's Guide*
- *IBM FAStT200 Fibre Channel Cabling Instructions*
- *IBM TotalStorage FAStT EXP700 Storage Expansion Unit Installation and User's Guide*
- *IBM FAStT EXP500 Installation and User's Guide*
- *IBM TotalStorage FAStT C2–133 Host Bus Adapter Installation and User's Guide*
- *IBM FAStT FC–2 Host Bus Adapter Installation and User's Guide*
- *IBM FAStT Host Adapter Installation and User's Guide*
- *IBM FAStT Remote Mirror Option Installation and User's Guide*
- *IBM Fibre Channel Basic SAN Configuration Setup Guide*

- *IBM FAStT MSJ User's Guide*
- *IBM Fibre Channel Problem Determination Guide*
- *IBM Fibre Channel Hardware Maintenence Manual*
- *IBM FAStT Storage Manager Script Commands*

You can also order publications through the IBM Publications Ordering System at the following Web site:

www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi/

Online help

Storage Manager 8.21 provides online help for the Enterprise Management and Subsystem Management windows. These help systems contain information about working with the management domain and about managing storage subsystems.

You can access the help systems from the Enterprise Management and Subsystem Management windows in Storage Manager 8.21. Click **Help** on the toolbar or press F1.

The help systems contain operating information that is common to all operating environments. For operating-system-specific information, see the following chapters and appendices of this document.
- For AIX, see Chapter 2, "Installing storage management station software on AIX systems" on page 11 and Appendix A, "AIX system requirements" on page 69.
- For HP-UX, see Chapter 3, "Installing storage management station software on HP-UX systems" on page 29 and Appendix B, "HP-UX system requirements" on page 71.
- For Solaris, see Chapter 4, "Installing storage management station software on Solaris systems" on page 43 and Appendix C, "Solaris system requirements" on page 73.

Web sites

For the most up-to-date information about IBM FAStT storage servers, go to the following Web site:

www.storage.ibm.com/hardsoft/disk/fastt/

For information about all IBM storage products, go to the following Web site:

www.ibm.com/storage/

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this document, you can submit them in one of the following ways:
- E-mail

  Submit your comments electronically to:

  starpubs@us.ibm.com

  Be sure to include the name and order number of the document and, if applicable, the specific location of the text you are commenting on, such as a page number or table number.

- Mail or fax

  Fill out the Readers' Comments form (RCF) at the back of this document and return it by mail or fax (1-800-426-6209) or give it to an IBM representative. If the RCF has been removed, you can address your comments to:

  International Business Machines Corporation
  RCF Processing Department
  Dept. M86/Bldg. 050-3
  5600 Cottle Road
  San Jose, CA   95193-0001
  U.S.A.

Chapter 1. Introduction

The IBM FAStT Storage Manager Version 8.21 for UNIX and AIX environments, known as *Storage Manager 8.21*, is a set of client and host tools that allow you to manage IBM FAStT200, FAStT500, and FAStT700 storage subsystems from a storage management station.

You can install Storage Manager 8.21 on a *storage management station*, which is the system that is responsible for managing all, or a portion of, a storage network. The storage management station communicates with the network management agents that reside in the managed nodes using a network management protocol, such as Simple Network Management Protocol (SNMP). Storage management commands are sent to the storage subsystem controllers, where the controller firmware validates and runs the commands, and then returns status and configuration information to the client software.

Typically, a storage management station is a remote system, connected to an Ethernet network, that is used to manage one or more storage subsystems. A storage management station can also be a host that is connected to the storage subsystem with a fibre-channel I/O path; you use this same path to manage the attached storage subsystems. Even though you can install the storage management software on a host, the host still uses the Transmission Control Protocol/Internet Protocol (TCP/IP) to communicate with the host-agent. The agent communicates with the controllers over the fibre-channel connection through the access volume.

This document provides system administrators with information about installing, configuring, and working with Storage Manager 8.21 in UNIX and AIX environments. Before installing Storage Manager 8.21, consult the following documentation:

README.txt files
> Read these first. Text files containing the latest installation and user information about the storage management software and hardware components are located in each operating system subdirectory on the installation CD. They are also maintained on the following Web site:
>
> www.ibm.com/pc/support/

IBM FAStT Storage Manager Concepts Guide
> Use this reference document to become familiar with the terminology and the features of the Storage Manager 8.21 software. This document is available on the installation CD and maintained at the following Web site:
>
> www.ibm.com/pc/support/

For information about installing Storage Manager 8.21 software on AIX, HP-UX, or Solaris systems, refer to this document. When you have completed the entire installation process, refer to the following online help systems, which contain information that is common to all operating system environments. You can access the help systems from the Enterprise Management and Subsystem Management windows in Storage Manager 8.21 by clicking **Help** on the toolbar or pressing F1.

Enterprise Management Help window
> Use this online help system to learn more about working with the entire management domain.

Subsystem Management Help window
> Use this online help system to learn more about managing individual storage subsystems.

Storage Manager 8.21 features

Storage Manager 8.21 supports FAStT200, FAStT500, and FAStT700 storage servers, and provides the features that are described in this section.

Event monitoring

Event monitoring allows you to monitor storage subsystems, even when the Enterprise Management window is not open. A toolbar button and an option, **Tools —> Update monitor**, allow updates to the monitor if changes have occurred to your storage subsystem configuration.

Heterogeneous environment

Heterogeneous Environment, with the storage partitioning feature, allows you to connect hosts that are running different operating systems to a single storage subsystem.

Configuration replication

Configuration replication allows you to save the logical configuration of a storage subsystem, and then load it on an identical storage subsystem. This feature can be used to replicate a logical configuration from one storage subsystem to another, or to save a storage subsystem configuration for backup.

Auto Volume Transfer (AVT)/Auto Disk Transfer (ADT)

Auto Volume Transfer (AVT), also known as Auto Disk Transfer (ADT) and referred to in this document as AVT/ADT, is a failover method. By this method, the storage subsystem initiates the transfer of a volume, or volumes, when it detects a controller or path failure. If the host detects the failure, it automatically sends I/O to the remaining functional path, simultaneously initiating an AVT/ADT event. The alternative to this approach is the Redundant Dual Active Controller (RDAC) model. By this model, the host initiates an explicit failover for all path failure events, except a controller failure.

Support for AVT/ADT varies by operating system:

- AIX requires that AVT/ADT be disabled. If AIX is in a heterogeneous environment with Solaris then AVT/ADT must be disabled.
- HP-UX requires that AVT/ADT be enabled. If HP-UX is in a heterogeneous environment with Solaris, then AVT/ADT must be enabled.
- Solaris supports AVT/ADT in a homogeneous environment.

> **Note:** For best performance in a Solaris environment, disable AVT/ADT.

Command-line interface

The command-line interface is based on the script-engine commands found in the script editor and allows you to issue commands to the storage subsystems. See the SMcli.txt file on the installation CD or the command line interface (CLI) online help topic for usage information.

Controller diagnostics

You can access the controller diagnostics from the **Controller —> Run Disagnostics** menu in the storage management software. These diagnostics allow you to test the host-side and drive-side fibre-channel loops.

Access volume mapping

Access volume mapping allows you to assign and change the LUN assignment for the access volume. An access volume is required for

communicating in a host-agent-managed environment using the fibre-channel input/output (I/O) path, but is not required for a directly-managed storage subsystem that is connected through the Ethernet cable connection.

High-availability cluster services
A high-availability clustering system allows application services to continue when a hardware or software failure occurs. This system protects you from software failures as well as from the failure of a CPU, disk, or local area network (LAN) component. If a component fails, the redundant component takes over cluster services and coordinates the transfer between components.

Redundant Dual Active Controller (RDAC)
The Redundant Dual Active Controller (RDAC) package, also known as the Redundant Disk Array Controller (RDAC) package, is a multipath device driver that provides controller failover support when a failure occurs anywhere along the fibre-channel I/O path.

Direct-attached configuration
Storage Manager 8.21 supports FAStT200, FAStT500, and FAStT700 storage servers in direct-attached homogenous environments.

Storage Manager 8 software packages

Storage Manager 8.21 contains the following software packages:
- Storage Manager 8 client software:
  - JRE (SMruntime)
  - SMclient
- Storage Manager 8 host software:
  - Storage Manager 8 agent (SMagent)
  - Storage Manager 8 utility package (SMutil)
  - Storage Manager 8 multipath device drivers (RDAC)

Install the software packages in the sequences shown in Table 1. For installation instructions, see the referenced sections. For an overview of each package, see the subsections that follow the table.

Table 1. Installation sequences of Storage Manager 8.21 software packages by host type

Step	AIX	HP-UX	Solaris
1	SMclient (see "Installing the client software on AIX hosts" on page 12)	SMruntime (see "Installing the host software on HP-UX hosts" on page 30)	SMruntime (see "Installing host software on Solaris hosts" on page 46)
2	RDAC (see "Installing host software on AIX hosts" on page 13)	SMclient (see "Installing the host software on HP-UX hosts" on page 30)	SMclient (on Solaris 8 only; see "Installing the client software on Solaris hosts" on page 45)
3		SMagent (see "Installing the SMagent on HP-UX hosts" on page 32)	RDAC (see "Installing host software on Solaris hosts" on page 46)

Step	AIX	HP-UX	Solaris
4		SMutil (see "Installing the SMutils on HP-UX hosts" on page 32)	SMagent (see "Installing host software on Solaris hosts" on page 46)
5			SMutil (see "Installing host software on Solaris hosts" on page 46)

Storage Manager 8 client software package

The Storage Manager 8.21 client software, SMclient, is a Java™-based GUI interface that allows you to configure, manage, and troubleshoot FAStT200, FAStT500, or FAStT700 storage servers as well as the EXP700 expansion enclosure in a storage subsystem through a host system or through a storage management station. Specifically, SMclient allows you to:

- Configure disk arrays and logical volumes
- Assign names to arrays and volume groups
- Assign logical volumes to storage partitions
- Replace and rebuild failed disk drives
- Expand the size of arrays
- Change RAID-type arrays
- Monitor the status of FAStT storage servers
- Perform troubleshooting and maintenance tasks, such as downloading firmware to controllers and upgrading NVSRAM
- Expand capacity using the EXP700 1 Gb or 2 Gb fibre-channel enclosure

Additionally, SMclient with FAStT700 allows you to:

- Loop four fibre-channel hosts using eight host connections with either 1 Gb or 2 Gb channels
- Assign redundant RAID controllers

Note: SMclient supports Solaris 2.6, 2.7, and 8.

The SMclient contains two main components:

Enterprise Management
> This component allows you to add, remove, and monitor storage subsystems in the management domain.

Subsystem Management
> This component allows you to manage the components of an individual storage subsystem.

For more information about this software, see either Chapter 2, "Installing storage management station software on AIX systems" on page 11, Chapter 3, "Installing storage management station software on HP-UX systems" on page 29, or Chapter 4, "Installing storage management station software on Solaris systems" on page 43 and the *IBM FAStT Storage Manager Concepts Guide.*

Storage Manager 8 agent software package

The Storage Manager 8.21 agent (SMagent) package contains the host-agent software, which you can use on HP-UX host systems to manage storage subsystems through the host fibre-channel connection. The host-agent software takes requests from a storage management station that is connected to the host through a network connection and passes the requests to the storage subsystem controllers through the fibre-channel I/O path.

For more information about managing storage subsystems through the host agent, see "Host-agent (in-band) management method".

Storage Manager 8 utility software package

Use the Storage Manager 8.0 utility (SMutil) package to register and map new logical drives to the operating system. Install SMutil on all HP-UX and Solaris host systems attached to a storage subsystem. The host computers are attached to the storage subsystem through the fibre channel.

RDAC

AIX and Solaris host systems require an RDAC driver for fibre-channel path redundancy. If a FAStT storage server has two controllers, and the operating system does not support multipath I/O, then you can use the RDAC. The RDAC monitors I/O paths; if a component failure occurs in one of the fibre-channel paths, the RDAC reroutes all I/O to another path.

Note: The AIX RDAC driver files are not included on the Storage Manager 8.21 installation CD; you must follow the instructions in "Installing host software on AIX hosts" on page 13 to download them from the appropriate Web site.

Storage subsystem management methods

The storage management software provides two methods for managing storage subsystems:

- The host-agent (in-band) management method. In this method, you manage the storage subsystems through the fibre-channel I/O path to the host.

  Note: You cannot use this management method on AIX systems.
- The direct (out-of-band) management method. In this method, you manage the storage subsystems directly over the network through the Ethernet connection to each controller.

Host-agent (in-band) management method

When you use the host-agent (in-band) management method, you manage the storage subsystems through the fibre-channel I/O path to the host. The management information can be processed by the host or passed to the storage management station through the network connection. Figure 1 on page 6 shows the host-agent (in-band) management method.

Note: You cannot use the host-agent (in-band) management method on AIX systems.

Storage subsystems

Controller

Controller

Running the
host-agent software

Controller

Controller

Host computer

Storage subsystems

Fibre Channel
I/O path

Network

Note: The host can also act as a
management station.

Management station
(one or more)

SJ000707

Figure 1. Host-agent (in-band) managed storage subsystems

Managing storage subsystems using the host-agent (in-band) management method
has the following advantages:

- You do not need to run Ethernet cables to the controllers.
- You do not need a Dynamic Host Configuration Protocol (DHCP) bootstrap
  protocol (BOOTP) server to connect the storage subsystems to the network.
- You do not need to configure the controller network (described in Chapter 2,
  "Installing storage management station software on AIX systems" on page 11,
  Chapter 3, "Installing storage management station software on HP-UX systems"
  on page 29, or Chapter 4, "Installing storage management station software on
  Solaris systems" on page 43).
- When adding devices, you need to specify a host name or Internet Protocol (IP)
  address for the host only, not for the individual controllers in a storage
  subsystem. Storage subsystems that are attached to the host are automatically
  discovered.

Managing storage subsystems using the host-agent (in-band) management method
has the following disadvantages:

- You are limited to configuring one less LUN than the maximum number allowed
  by the operating system and host adapter that you are using.
- The host-agent requires a special logical drive, called an *access volume*, to
  communicate with the controllers in the storage subsystem.

- If you are upgrading controllers from firmware version 3.x to version 4.x and your host system has already configured its maximum number of LUNs, you must give up a LUN to be used as an access volume.

Important: The access volume uses one of the LUNs. If your host already has the maximum number of LUNs configured, either use the direct-management method or give up a LUN for use as the access volume. For information about your specific configuration, see the appropriate chapter in this document for your operating system environment.

Direct (out-of-band) management method

When you use the direct (out-of-band) management method, you manage storage subsystems directly over the network through the Ethernet connection to each controller. To manage the storage subsystem through the Ethernet connections, you must define the IP address and host name for each controller and attach a cable to the Ethernet ports on each of the storage subsystem controllers. Figure 2 shows the direct (out-of-band) management method.

Note: You can use the direct (out-of-band) management method for AIX, HP-UX, and Solaris systems.
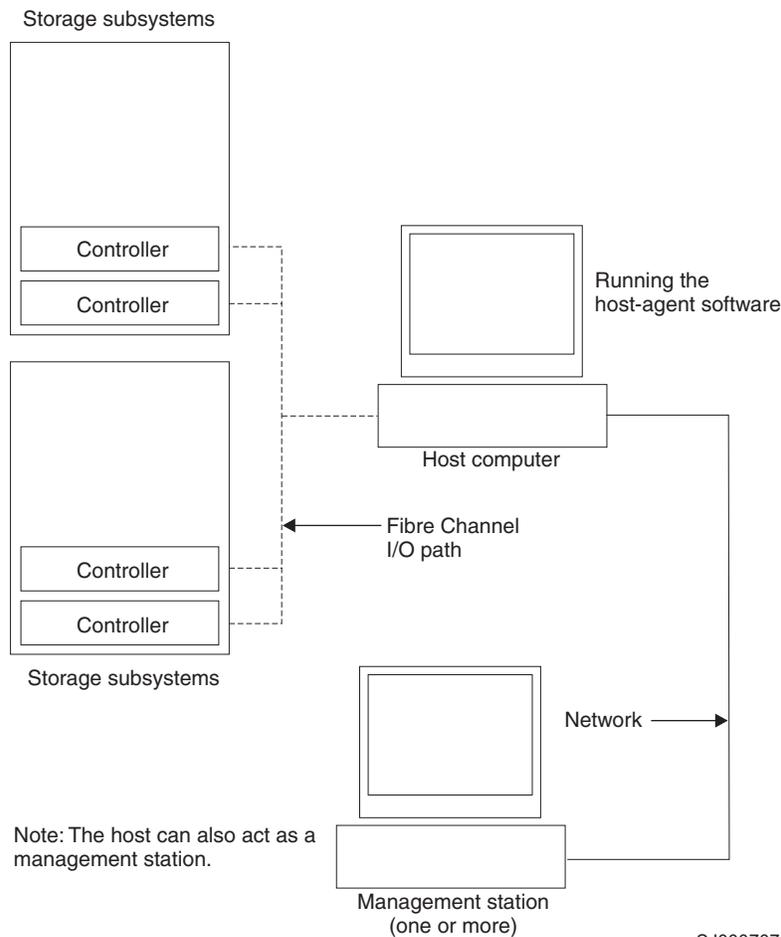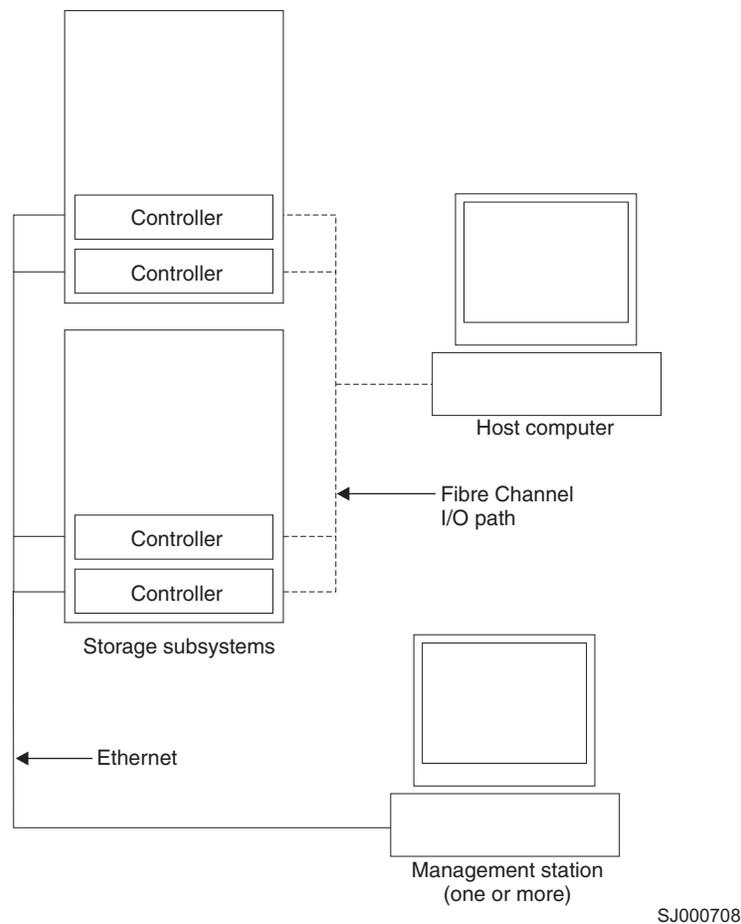

Figure 2. Direct (out-of-band) managed storage subsystems

Managing storage subsystems using the direct (out-of-band) management method has the following advantages:

- The Ethernet connections to the controllers enable a storage management station running SMclient to manage storage subsystems that are connected to a host running one of the following operating systems:
  – AIX
  – HP-UX
  – Linux
  – Microsoft Windows NT/2000
  – Solaris
  – Any other operating systems that are supported by Storage Manager 8.21
- For AIX and HP-UX host systems, you do not need to use an access LUN to communicate with the controllers. You can configure the maximum number of LUNs that are supported by the operating system and the host adapter that you are using.

  Note: Solaris host systems always use access LUNs.

Managing storage subsystems using the direct (out-of-band) management method has the following disadvantages:
- It requires two Ethernet cables to connect both storage subsystem controllers to the network.
- When you configure IP addresses for new devices, you must either assign a static IP address or host name for each controller, or place the devices on a network with a DHCP or BOOTP server.
- Other network preparation tasks might be required. For more information, see the system installation guide for your network server.

Operating system requirements

Storage Manager 8.21 supports the FAStT200, FAStT500, and FAStT700 storage servers in the following environments:
- IBM AIX 4.3.3 or 5.1 (see Appendix A, "AIX system requirements" on page 69)
- HP-UX 11.0 (32 bit or 64 bit), HP-UX 11.11, or HP-UX 11i (see Appendix B, "HP-UX system requirements" on page 71)
- Sun Solaris 2.6, 7, or 8 (see Appendix C, "Solaris system requirements" on page 73)

Setting up IP addresses for FAStT storage controllers

Contact your network administrator to obtain the IP address and associated host name for each controller in every storage subsystem on the network, and make a note of those values for reference.

To use the out-of-band management method without setting up a DHCP BOOTP server, you must assign IP addresses to the FAStT controllers using CLI commands through serial cables that are connected to a terminal emulator.

Complete the following steps to set up the FAStT controller IP addresses using serial ports:
1. Stop all I/O to the FAStT controllers.
2. Connect a serial cable from one of the controllers to a system with a terminal emulator available.

Note: Use HyperTerminal Version 3 (V3) or later.

3. Connect to the FAStT storage server and send a break signal (Ctrl+Break for most emulators).

4. Repeat this step until the following message is displayed:

```
Press the space bar for baud rate within 5 seconds.
```

5. Press the space bar to ensure the correct baud rate setting.

6. Send another break signal; the following message is displayed:

```
Press within 5 seconds: ESC for SHELL, BREAK for baud rate.
```

7. Press Escape to access the shell of the controller.

8. Type the password: `infiniti`

9. Type `netCfgShow` to show the current network configuration.

10. Type `netCfgSet` to change the network configuration information.

11. Assign an IP address to the controller.

12. Disconnect from the first controller and connect to the second controller.

13. Repeat steps 1 on page 8 - 12 to assign the second IP address to the second controller.

14. Restart the FAStT storage server by turning the controller unit on and off.

Chapter 2. Installing storage management station software on AIX systems

This chapter provides the following specific information for AIX operating systems:

- Hardware and firmware requirements
- Client software installation
- Host software installation
- AIX restrictions
- AIX configuration information

Hardware and firmware requirements

Table 2 lists the supported versions of hardware to use with Storage Manager 8.21.

Table 2. Supported versions of hardware for AIX systems

Product Name	Model	Product release and firmware version
IBM FAStT700 RAID Controller Enclosure Unit	1742-1RU	Firmware 05.21.25.00, NVSRAM CNV1742R821NT008
IBM FAStT700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9319
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Firmware 05.21.25.00, NVSRAM CNV3352R821NT008
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9163, 9165
IBM FAStT200 RAID Controller Enclosure Unit	3542-1RU and FC 2101	Firmware 05.20.12.00, NVSRAM CNV3542R821NT018
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Firmware 05.20.12.00, NVSRAM CNV3542R821NT018
IBM HA Emulex LP7000	FC 6227	3.22A1
IBM HA Emulex LP9000	FC 6228	3.82A1
Brocade switch	FC 2109	2.6.0.c/3.0.2.c
McData switch	IBM 2031-16	02.00.00.25
McData switch	IBM 2032-001	02.00.00.25
McData switch	IBM 2032-064	02.00.00.25
INRANGE switch	IBM 2042	3.2.1.J

Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700

Installing the client software on AIX hosts

Use the following procedure to install the client software on an AIX storage management station or on an AIX host acting as a storage management station.

Prerequisites

Before installing the client software, ensure that the following conditions are met:

- The AIX host on which you are installing the SMclient software meets the minimum hardware and software requirements described in "Hardware and firmware requirements" on page 11 and "Software requirements" on page 69.
- The correct filesets are present on the system.

  Note: If the filesets are not present, follow the instructions in "Installing host software on AIX hosts" on page 13 to download them from the appropriate Web site.

  – For a list of AIX 4.3.3 filesets, see Table 11 on page 69.
  – For a list of AIX 5.1 filesets, see Table 12 on page 69.

Installing the SMclient on AIX hosts

In the following procedure, the installation CD is mounted at /dev/cdrom. Adjust these instructions as required for your specific installation. No restart is required during the client installation process.

1. Create the /cdrom mount point:

```
# mkdir /cdrom
```

2. Insert the installation CD into the CD-ROM drive and mount the CD at /cdrom:

```
# mount -v cdrfs -o ro /dev/cd0 /cdrom
```

3. Install the SMclient software package:

```
# installp -a -d /cdrom/AIX/SMclient.aix-08.21.G5.02.bff SMclient.aix.rte
```

   The installation process begins. The process displays information as it runs, including an installation summary when the process is finished.

4. Verify that the installation was successful:

```
# lslpp -ah SMclient.aix.rte
```

   The verification process returns a table that describes the software installation, including the install package file name, version number, action, and action status. If the verification process returns an error, contact your customer support representative.

5. You are finished with the client software installation for this AIX storage management station. Unmount the CD using the following command, and then remove it from the CD-ROM drive:

```
# umount /cdrom
```

Installing host software on AIX hosts

After you install the Storage Manager 8.21 client software and configure your storage subsystems, use these instructions to install the appropriate Storage Manager 8.21 host software. See Table 1 on page 3 for information about the software packages and the installation sequence required by each host type.

All AIX hosts in your storage subsystem must have the RDAC multipath driver installed on them. This section describes how to install that host software package.

Prerequisites

This section lists the prerequisites for installing the RDAC driver program temporary fixes (PTFs) on an AIX system and describes how to verify AIX host firmware.

Prerequisites for installing RDAC

See Table 11 on page 69 or Table 12 on page 69 for a list of PTF filesets required by version.

Verifying AIX host firmware

Use the **lslpp** command to verify that the correct version numbers of the AIX host software are installed:

```
# lslpp -ah filename
```

Where *filename* is one of the required filesets listed in Table 11 on page 69 or Table 12 on page 69.

For example:

```
# lslpp -ah devices.fcp.disk.array.rte
```

Installing RDAC on AIX hosts

Complete the following procedure to update the RDAC driver PTF (devices.fcp.disk.array.rte) on an AIX system. Repeat this procedure for all AIX systems that are connected to the storage subsystem.

1. Go to one of the following Web sites:

    ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
    ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
    ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700

2. Click **Downloads** in the downloads section.
3. Scroll to the operating system-specific updates section of the Web page.
4. Follow the link to the appropriate sets of files for your system and follow the installation instructions.
5. Verify that the correct version of the software was successfully installed:

```
# lslpp -ah devices.fcp.disk.array.rte
```

    The verification process returns a table that describes the software installation, including the installation package fileset name, version number, action, and action status. If the verification process returns an error, contact your customer support representative. If it does not return an error, then you are finished installing the RDAC driver on this AIX system.

6. Shut down and restart the AIX system. This must be done for the software changes to take effect.

```
# shutdown -Fr
```

Performing the initial configuration of storage subsystems on AIX hosts

All FAStT storage servers ship with NVSRAM configured for Windows NT hosts and have an access volume set to LUN 31. However, because AIX does not support access volumes, you must perform the following tasks which are described in detail (along with the other configuration tasks) starting with 1.

- Disable AVT/ADT using the script that is supplied on the Storage Manager 8.21 installation CD-ROM. Follow the instructions in 9 on page 16. AVT/ADT must be disabled for path redundancy to be effective.

- Delete access volumes (also known as UTM LUNs). Follow the instructions in 6 on page 15. Access volumes are not supported; failure to delete them can cause system errors.

- If the controllers are shipped with a firmware release prior to version 05.20.08.00 and an NVSRAM release prior to version CNV3542R821NT016 (for FAStT200 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV3552R821NT007 (for FAStT500 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV1742R821NT007 (for FAStT700 storage servers), they must be upgraded. Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

  ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
  ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
  ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700

Note: The versions of firmware and NVSRAM that ship with this product can be later releases than those described in this document.

Complete the following procedure to configure Storage Manager 8.21 for an AIX system.

1. To set up the storage subsystem for AIX, the subsystem must be physically configured for direct management through the Ethernet connections on each controller. Install SMclient before configuring the subsystem.

   Note: See "Setting up IP addresses for FAStT storage controllers" on page 8 for information about assigning IP addresses to the controllers.

2. After the disk subsystem is configured on the network, start the SMclient software on the host server:

```
# /usr/SMclient/SMclient
```

3. Complete the following steps to specify the IP addresses of the controllers.
   a. In the Enterprise Management window, click **Edit —> Add Device**.
   b. In the Add Device window, type the IP address of the first controller in the storage subsystem and click **Add**.

c. Type the IP address of the second controller and click **Add**, and then click **Close**.

The storage subsystem is shown as a direct network attachment. Double-click **Storage Subsystem** to open the Subsystem Management window.

4. In the Subsystem Management window, click **View—> Storage Subsystem Profile** and review the summary portion of the output. Verify that the controller firmware and NVSRAM are at the correct versions. If they are, go to step 10 on page 16; if not, continue with step 5.

   Note: Upgrade the NVSRAM, and then ensure that the controllers are in a quiescent state before you upgrade the controller firmware. Concurrent upgrades of NVSRAM are not supported.

5. Complete the following steps to upgrade the NVSRAM.

   a. Mount and locate the NVSRAM file on the installation CD (`/cdrom/NVSRAM/3542`, `/cdrom/NVSRAM/3552`, or `/cdrom/NVSRAM/1742`), or download the correct version from the Web site. Place the file in a designated directory on the host system.

   b. In the Subsystem Management window, click **Storage Subsystem —> Download —> NVSRAM**.

   c. Type or select the full pathname of the NVSRAM directory.

   d. Double-click the NVSRAM filename or click **OK** to select the correct file.

   e. Click **Yes** to start the download.

   f. Verify that the NVSRAM was successfully installed.

6. Complete the following steps to delete access volumes.

   a. In the Subsystem Management window, click the **Logical Mapping** tab.

   b. Right-click **Access Logical Drive**.

   c. Click **Delete**.

   d. Type `yes`.

7. After you ensure that all I/O to the controllers is stopped, complete the following steps to upgrade the controller firmware.

   a. Mount and locate the firmware file on the installation CD (`/cdrom/Firmware/3542`, `/cdrom/Firmware/3552` or `/cdrom/Firmware/1742`), or download the correct version from the Web site. Place the file in a designated directory on the host system.

   b. In the Subsystem Management window, click **Storage Subsystem —> Download —> Firmware**.

   c. Type or select the full pathname of the firmware file.

   d. Click **OK** to update both controllers.

   e. Click **Yes** to start the download. A new window opens.

   f. Close the Subsystem Management window and then reopen it to complete the firmware update.

   g. Verify that the firmware was successfully installed.

8. After you upgrade the NVSRAM and the firmware, turn the power to the controller enclosure off and on in the following sequence, and then restart the AIX host.

   a. Turn off both controllers.

   b. Wait at least 30 seconds.

   c. Turn on both controllers.

9. Complete the following steps to disable AVT/ADT. For more information about AVT/ADT, see Chapter 5, "Completing the software installation" on page 55.

   a. Insert the installation CD into the CD-ROM drive of the Storage Management Station.

   b. Start the storage management software.

   c. In the Enterprise Management window, highlight the storage subsystem on which you want to disable AVT/ADT.

   d. Click **Tools —> Execute Script**. The script editor window opens.

   e. Click **File —> Load Script**.

   f. Select the `avtdisable.scr` file form the `/cdrom/scripts` directory of the installation CD.

   g. Click **OK** to load the script.

   h. Click **Tools —> Execute Only**. The script runs. It displays a script execution complete message when it is finished.

   i. Turn off the controller enclosure, and then turn it back on.

10. You must install the RDAC driver on the AIX system before you can use it to recognize and configure FAStT storage servers. See "Verifying AIX host firmware" on page 13 to verify that the driver is installed. If the RDAC driver is not installed, install it now and then use the following command to probe for devices:

```
# cfgmgr -v
```

> **Note:** For information about creating logical drives and partitions, see the *IBM FAStT Storage Manager Concepts Guide*, available on the installation CD or at the following Web site:
>
> www.ibm.com/pc/support/

11. Click the **Mappings View** tab on the Subsystem Management window.

12. In the Mappings window, create a new host group by clicking **Mappings —> Define —> Host Group**.

13. Type the name of the new host group (for example, AIX). Click **Add**, and then click **Close**.

14. Highlight the new host group and click **Mappings —> Define —> Host**.

15. Define the new host. Type the name of the AIX host to which the storage subsystem is attached.

   a. Click **Add**, and then click **Close**.

   b. Highlight the host that you just added and right-click **Define New Host Port**.

   c. Select the desired host port for the first HBA, and then change the host type to **AIX** and click **Add**.

   > **Note:** To verify that the host port matches the AIX host, run the **lsdev -Cc adapter | grep fcs** command. Identify the fcs by associating the number of HBAs, then run the **lscfg -vl fcs# |grep** Network command. Verify that the number that displays under Network Address matches the host port number in the GUI. Repeat this procedure for the second host port.

   d. Choose the host port for the second HBA and click **Add**, and then click **Close**.

AIX restrictions

The following restrictions apply to FAStT200, FAStT500, and FAStT700 storage servers:

- FC Boot Support is not available.
- F-RAID Manager is not supported.
- The maximum number of partitions per AIX host, per FAStT storage server, is 2.
- Each AIX host can support two or four host bus adapters (FC 6227 or 6228) and up to two FAStT storage partitions, each requiring two host bus adapter connections.
- Direct-attach configurations are restricted to single-initiator configurations only. In other words, only one connection to each FAStT Mini-hub is allowed.
- Single HBA configurations are not allowed. Each connection to a partition must have 2 host bus adapters configured.
- Single switch configurations are permitted, but each host bus adapter and FAStT Controller combination must be in a separate SAN zone.
- All volumes configured for AIX must be mapped to an AIX Host group. Connecting and configuring to volumes in the ″default host group″ is not permitted. (See "Performing the initial configuration of storage subsystems on AIX hosts" on page 14.)
- Other storage devices, such as tape devices or other disk storage, must be connected through separate host bus adapters and SAN zone.

AIX configuration information

This section contains the following AIX configuration information:

- Installing and setting attributes of the RDAC driver for AIX
- Identifying the controller ID numbers
- Identifying device names and bus numbers
- Identifying logical drives by operating system device names
- Redistributing LUNs in case of failure
- Creating a direct-attached configuration

Installing and setting attributes of the RDAC driver for AIX

The RDAC driver must be installed on all AIX hosts that will be attached to a FAStT storage subsystem. The RDAC driver creates the following devices that represent the FAStT storage subsystem configuration:

darX   The disk array router represents the entire array, including current and deferred paths to all LUNs (hdisks on AIX).

dac0, dac1
The disk array controller devices represent the two controllers of the array.

hdiskX
These devices represent individual LUNs on the array.

When these devices are configured, the Object Data Manager (ODM) is updated with default parameters. You must change some of these parameters for AIX to operate properly, while others can be changed to enhance the operability of your configuration. Use the **lsattr -El** command to view attribute settings on an AIX system, as shown in the following examples.

This example shows the attribute settings for a disk array router (dar0):

```
# lsattr -El dar0
act_controller  dac0,dac1                                Active Controllers                         False
all_controller  dac0,dac1                                Available Controllers                      False
held_in_reset   none                                     Held-in-reset controller                   True
load_balancing  no                                       Dynamic Load Balancing                     True
autorecovery    no                                       Autorecover after failure is corrected     True
lun_bitmap      0000000000000000000000000000001010       LUN Ownership bitmap                        True
hlthchk_freq    600                                      Health check frequency in seconds          True
aen_freq        600                                      Polled AE frequency in seconds             True
balance_freq    600                                      Dynamic Load Balancing frequency in seconds True
fast_write_ok   yes                                      Fast Write available                       False
cache_size      512                                      Cache size for both controllers            False
switch_retries  5                                        Number of times to retry failed switches   True
```

This example shows the attribute settings for dac0:

```
# lsattr -El dac0
passive_control no                Passive controller         False
alt_held_reset  no                Alternate held in reset    False
controller_SN   1T04611221        Controller serial number   False
ctrl_type       3552              Controller Type            False
cache_size      512               Cache Size in MBytes        False
scsi_id         0x210013          SCSI ID                    False
lun_id          0x0               Logical Unit Number         False
utm_lun_id      none              Logical Unit Number         False
location                          Location Label              True
ww_name         0x200800a0b80c1948 World Wide Name            False
GLM_type        low               GLM type                    False
```

This example shows the attribute settings for a LUN on the array (hdisk4):

```
# lsattr -El hdisk4
pvid          none                       Physical volume identifier              False
q_type        simple                     Queuing Type                            False
queue_depth   30                         Queue Depth                             True
reserve_lock  yes                        RESERVE device on open                  True
write_cache   yes                        Write Caching enabled                   True
size          69270                      Size in Mbytes                          False
raid_level    5                          RAID Level                              False
rw_timeout    30                         Read/Write Timeout value                True
reassign_to   120                        Reassign Timeout value                  True
scsi_id       0x210013                   SCSI ID                                 False
lun_id        0x0000000000000000         Logical Unit Number                     False
cache_method  fast_write                 Write Caching method                    True
prefetch_mult 0                          Multiple of blocks to prefetch on read  True
ieee_volname  600A0B80000C1A70000000013B45FEA8 IEEE Unique volume name           False
```

Required attribute settings

This section lists the required attribute settings for dar and hdisk devices and shows how to set them using the **chdev -l** command. To make the attribute changes permanent in the Customized Devices object class, use the **-P** option.

Attribute settings for dar devices: For multi-initiator configurations, the
autorecovery attribute must be set to **no**.

```
# chdev -l dar0 -a autorecovery=no
```

On single host systems, the *load_balancing* attribute can be set to **yes** to enhance
performance.

```
# chdev -l dar0 -a load_balancing=yes
```

Note: On multihost systems, the *load_balancing* attribute must be set to **no**.

Attribute settings for hdisk devices: Setting the *queue_depth* attribute to the
appropriate value is important for system performance. For large, multihost
configurations, always set the attribute to less than 10.

Use the following formula to determine the maximum queue depth for your system:

```
512 / (number-of-hosts * LUNs-per-host )
```

For example, a system with four hosts, each with 32 LUNs (the maximum number
of LUNs per AIX host), would have a maximum queue depth of 4:

```
512 / ( 4 * 32 ) = 4
```

In this case, you would set the *queue_depth* attribute for hdiskX as follows:

```
# chdev -l hdiskX -a queue_depth=4 -P
```

Attention: If you do not set the queue depth to the proper level, you might
experience loss of filesystems and system panics.

The *write_cache* attribute does not control whether or not write caching is on; the
FAStT controllers do. By default, write caching is on for the FAStT controllers.

Definitions of attribute settings
The following tables provide information about the attributes of the following device
types:
- dar
- dac
- LUN

Table 3. Attributes for dar devices

Attribute	Definition
act_controller	List of controllers in the active state at the time of configuration.
all_controller	List of controllers that comprise this array; usually there are two dac devices.
held_in_reset	Name of the controller that was in the held-in-reset state at the time of configuration, or **none** if no controllers were in that state.

Table 3. Attributes for dar devices (continued)

Attribute	Definition
load_balancing	Indicator that shows whether load balancing is enabled (**yes**) or disabled (**no**); see the definition of the *balance_freq* attribute for more information.
autorecovery	Indicator that shows whether the device returns the array to dual-active mode when it detects proper operation of both paths and controllers (**yes**) or not (**no**).
lun_bitmap	Bitmap that shows which controller own which LUN at the time of configuration.
hlthchk_freq	Number that specifies how often health checks are performed, in seconds.
aen_freq	Number that specifies how often polled AEN checks are performed, in seconds.
balance_freq	If *load_balancing* is enabled, number that specifies how often the system performs load-balancing on the array, in seconds.
fast_write_ok	Indicator that shows whether fast-write write-caching is available for this system (**yes**) or not (**no**).
cache_size	Cache size for both controllers, in megabytes; **0** if the sizes do not match.
switch_retries	Number that specifies how many times to retry failed switches, in integers.

Table 4. Attributes for dac devices

Attribute	Definition
passive_control	Indicator that shows whether this controller was in passive state at the time of configuration (**yes**) or not (**no**).
alt_held_reset	Indicator that shows whether the alternate controller was in the held-in-reset state at the time of configuration (**yes**) or not (**no**).
controller_SN	Serial number of this controller.
ctrl_type	Type of array this controller belongs to; a value of **3542** indicates FAStT200; a value of **3552** indicates FAStT500; a value of **1742** indicates FAStT700.
cache_size	Cache size of this controller, in megabytes.
scsi_id	SCSI identifier of this controller.
lun_id	Logical unit number of this controller.
utm_lun_id	Logical unit number of this controller, or **none** if UTM (access volumes) is not enabled.
location	User-defined location label for this controller; the system does not use this value.
ww_name	Fibre-channel worldwide name of this controller.
GLM_type	GLM type used for this controller.

Table 5. Attributes for LUN devices

Attribute	Definition
pvid	AIX physical volume identifier, or **none** if not set.
q_type	Queueing type for this device; must be set to **simple**.
queue_depth	Number that specifies the depth of the queue based on system configuration; reduce this number if the array is returning a BUSY status on a consistent basis.
reserve_lock	Indicator that shows whether the fcparray issues a SCSI Remove command every time a device is opened or when a Test Unit Ready sequence is issued by the driver (**yes**) or not (**no**).
write_cache	Indicator that shows whether write-caching is enabled on this device (**yes**) or not (**no**); see the definition of the *cache_method* attribute for more information.
size	Size of this LUN.
raid_level	Number that specifies the RAID level of this device.
rw_timeout	Number that specifies the read/write timeout value for each read/write command to this array, in seconds; usually set to 30.

Table 5. Attributes for LUN devices (continued)

Attribute	Definition
reassign_to	Number that specifies the timeout value for FC reassign operations, in seconds; usually set to 120.
scsi_id	SCSI identifier at the time of configuration.
lun_id	Logical unit number of this device.
cache_method	If *write_cache* is enabled, the write-caching method of this array; set to one of the following: • **default**. Default mode; the word ″default″ is not seen if *write_cache* is set to yes. • **fast_write**. Fast-write (battery-backed, mirrored write-cache) mode. • **fw_unavail**. Fast-write mode was specified but could not be enabled; write-caching is not in use. • **fast_load**. Fast-load (non-battery-backed, non-mirrored write-cache) mode. • **fl_unavail**. Fast-load mode was specified but could not be enabled.
prefetch_mult	Number of blocks to be prefetched into read cache for each block read.
ieee_volname	IEEE unique volume name identifier for this LUN.

Initial device identification

After the FAStT storage subsystem has been set up, volumes have been assigned to the host, and the RDAC driver has been installed, use the following command to probe for the new devices:

```
# cfgmgr -v
```

Next, use the **lsdev -Cc disk** command to see if the device driver recognizes each FAStT200 LUN as a "3542 Disk Array Device", each FAStT500 LUN as a "3552 Disk Array Device" or each FAStT700 LUN as a "1742 Disk Array Device". The following example shows the results of the command for a set of FAStT500 LUNs:

```
# lsdev -Cc disk
hdisk0 Available 10-88-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 20-58-01    3552 Disk Array Device
hdisk2 Available 20-58-01    3552 Disk Array Device
hdisk3 Available 20-58-01    3552 Disk Array Device
hdisk4 Available 20-58-01    3552 Disk Array Device
```
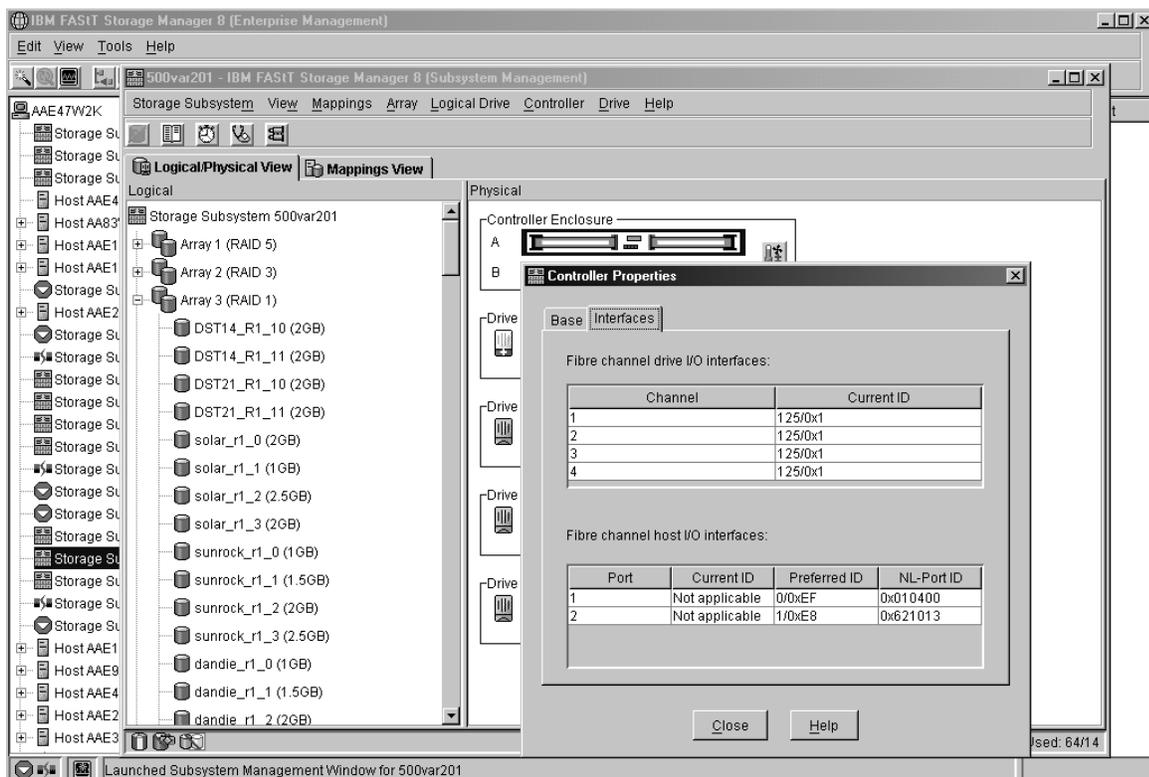
Identifying the controller ID numbers

One of the attributes listed by the **lsattr -El** command is the controller serial number (*controller_SN*) of that dac.

To display the FAStT Controller Properties window:

1. In the Physical view of the Subsystem Management window, select a controller.
2. Right-click **Controller —> Properties**.

   Figure 3 shows the Controller Properties window of the controller in slot A. This controller has a serial number of 1T04810361 and is represented by dac1.

```
# lsattr -El dac1
passive_control  no                    Passive controller      False
alt_held_reset   no                    Alternate held in reset False
controller_SN    1T04810361            Controller serial number False
ctrl_type        3552                  Controller Type         False
cache_size       512                   Cache Size in MBytes    False
scsi_id          0x210513              SCSI ID                 False
lun_id           0x0                   Logical Unit Number     False
utm_lun_id       none                  Logical Unit Number     False
location                               Location Label          True
ww_name          0x201200a0b807b856    World Wide Name         False
GLM_type         low                   GLM type                False
```


SJ000709

Figure 3. Controller Properties window

Identifying device names and bus numbers

After the operating system device names are found, those names must be correlated to the primary and secondary paths of the device, and then from each path to its associated logical drive.

Use the **lsdev**, **fget_config**, and **lsattr** commands to get information about device names and bus numbers.

Using the lsdev command

This section shows how to use the **lsdev** command to get information about disk array routers (dars), disk array controllers (dacs), and hdisks.

Note: The examples in this section show dars of type 3552 in a FAStT500 storage subsystem. The dars would be of type 3542 in a FAStT200 storage subsystem or type 1742 in a FAStT700 storage subsystem.

The following example uses the **lsdev** command to show the status of the dar, which represents the entire FAStT array. "darX" keeps track of the status of each dacX, and of the current and preferred paths to each hdisk.

```
# lsdev -C |grep dar
dar0      Available               3552 Disk Array Router
```

The following example uses the **lsdev** command to show the two dacs that represent the disk array controllers. The third column is the location code column. Two distinct paths are represented by the values 91-08-01 and 11-08-01. Each AIX system has its own set of location codes that describe the internal path of that device, including bus and host-adapter locations. See the service manual for your system type to identify device locations.

```
# lsdev -C |grep dac
dac0      Available 91-08-01     3552 Disk Array Controller
dac1      Available 11-08-01     3552 Disk Array Controller
```

The following example uses the **lsdev** command to show that there are eight disks (LUNs) of type 3552 recognized by this AIX host system (this is for a FAStT500; disks of type 3542 would be found for a FAStT200, and disks of type 1742 would be found for a FAStT700). The third column shows the location code of each device. Notice that there are four disks per path, or four disks per dac (controller).

```
# lsdev -Cc disk
hdisk0 Available 40-60-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 91-08-01     3552 Disk Array Device
hdisk2 Available 91-08-01     3552 Disk Array Device
hdisk3 Available 91-08-01     3552 Disk Array Device
hdisk4 Available 11-08-01     3552 Disk Array Device
hdisk5 Available 11-08-01     3552 Disk Array Device
hdisk6 Available 11-08-01     3552 Disk Array Device
hdisk7 Available 11-08-01     3552 Disk Array Device
hdisk8 Available 91-08-01     3552 Disk Array Device
```

Using the fget_config command

The **fget_config** command displays the current owner of each hdisk. The following example shows that it is a quick way to determine which LUN (hdisk) is actively owned by a controller.

```
# fget_config -l dar0
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk1
dac0-hdisk2
dac0-hdisk3
dac1-hdisk4
dac1-hdisk5
dac1-hdisk6
dac1-hdisk7
dac0-hdisk8
```

Using the lsattr command

You can also use the **lsattr** command to find out which controller owns a particular hdisk. The following example shows portions of the outputs from several **lsattr** commands.

```
# lsattr -El dac0
passive_control  no                   Passive controller       False
alt_held_reset   no                   Alternate held in reset  False
controller_SN    1T01710113           Controller serial number False
ctrl_type        3552                 Controller Type          False
cache_size       512                  Cache Size in MBytes      False
scsi_id          0x11100              SCSI ID                  False
lun_id           0x0                  Logical Unit Number      False


# lsattr -El dac1
passive_control  no                   Passive controller       False
alt_held_reset   no                   Alternate held in reset  False
controller_SN    1T03910039           Controller serial number False
ctrl_type        3552                 Controller Type          False
cache_size       512                  Cache Size in MBytes      False
scsi_id          0x11000              SCSI ID                  False
lun_id           0x0                  Logical Unit Number      False


# lsattr -El hdisk1
----------- Parts removed:
scsi_id          0x11100              SCSI ID                  False
lun_id           0x0000000000000000   Logical Unit Number      False
cache_method     fast_write           Write Caching method     True
```

In the example, hdisk1 belongs to the controller represented by `dac0`. The hdisk has the same SCSI ID as the controller (dac) that owns it.
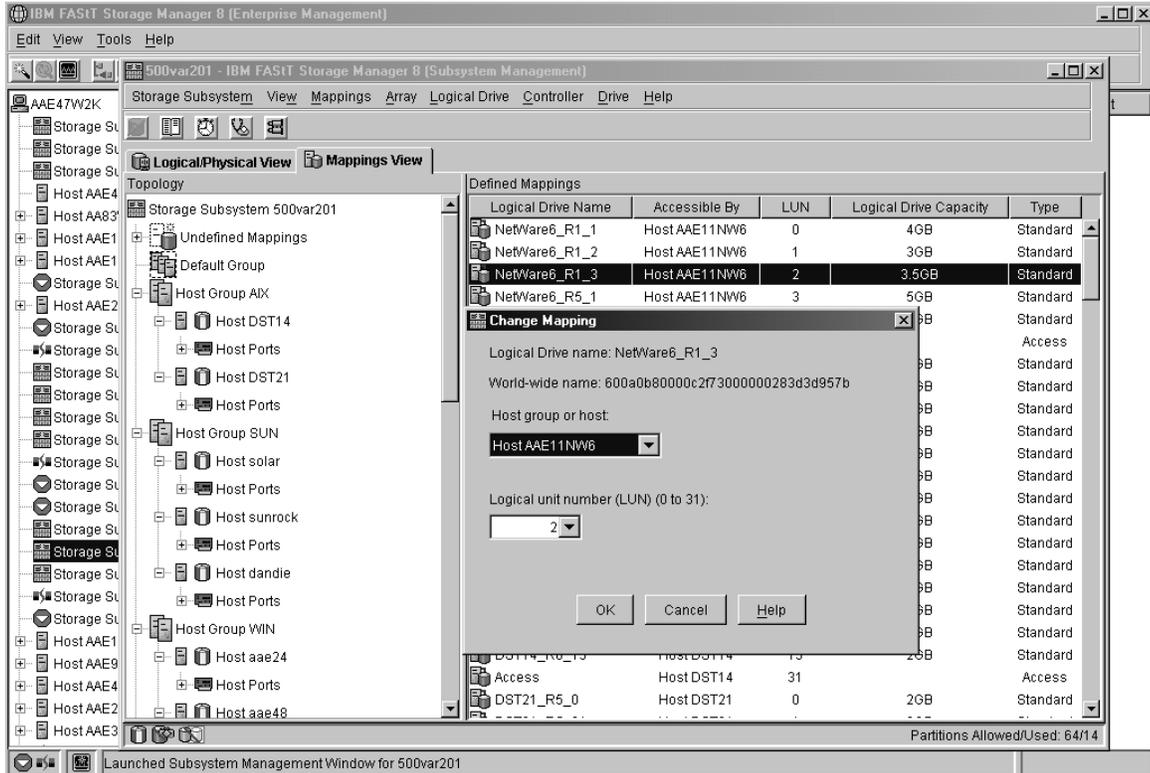
Identifying logical drives by operating system device names

The **lsattr** command provides detailed information about a LUN, including information that allows you to map the system device name to the logical volume on the FAStT storage subsystem.

For example, if you use the **lsattr** command on the LUN named hdisk4, it provides the following information: it is a 6 GB LUN of type RAID, with a LUN ID of 2, and an IEEE volume name of 600A0B80000C1E6300000033BB39927 (see the following example). You can make a quick identification by locating the LUN ID on the far right side of the **Mappings View**tab.

```
# lsattr -El hdisk4
pvid          none                      Physical volume identifier        False
q_type        simple                    Queuing Type                      False
queue_depth   30                        Queue Depth                       True
reserve_lock  yes                       RESERVE device on open            True
write_cache   yes                       Write Caching enabled             True
size          6062                      Size in Mbytes                    False
raid_level    0                         RAID Level                        False
rw_timeout    30                        Read/Write Timeout value          True
reassign_to   120                       Reassign Timeout value            True
scsi_id       0x11000                   SCSI ID                           False
lun_id        0x0002000000000000        Logical Unit Number               False
cache_method  fast_write                Write Caching method              True
prefetch_mult 0                         Multiple of blocks to prefetch on read True
ieee_volname  600A0B80000C1E63000000033BB39927 IEEE Unique volume name    False
```

You can make a more exact correlation using the distinctive *ieee_volname* attribute. The value of this attribute on the AIX host is the same as the Unique Logical Drive Identifier on the FAStT storage subsystem. The **ieee_volname** and Unique Logical Drive Identifier can be found in the Change Mappings window, which is shown in Figure 4. (Click **Mappings view** tab. Select **Logical Drive Name**. Right click and select **Change Mappings**.)

SJ000712

Figure 4. Change Mappings window

Another way to identify the system disk name based on the logical volume name is by using the Logical Drive Property window. Open this window by clicking the

Storage Management window and right-clicking **Logical Drive Name —>
Properties**. Look for the **World-wide name**, **Capacity**, and **RAID level** properties,
as shown in Figure 5.

The following example shows a portion of the output of the **lsattr** command for
hdisk20. Notice that the *ieee_volname* and *lun_id* attributes have the same values
as the **World-wide name** and **Subsystem ID (SSID)** properties shown in Figure 5,
respectively.

```
# lsattr -El hdisk20
write_cache   yes                           Write Caching enabled                     True
size          9216                          Size in Mbytes                            False
raid_level    5                             RAID Level                                False
scsi_id       0x210513                      SCSI ID                                   False
lun_id        0x0012000000000000            Logical Unit Number                       False
cache_method  fast_write                    Write Caching method                      True
prefetch_mult 0                             Multiple of blocks to prefetch on read    True
ieee_volname  600A0B80000C1E63000000163BB39B7C IEEE Unique volume name                False
```
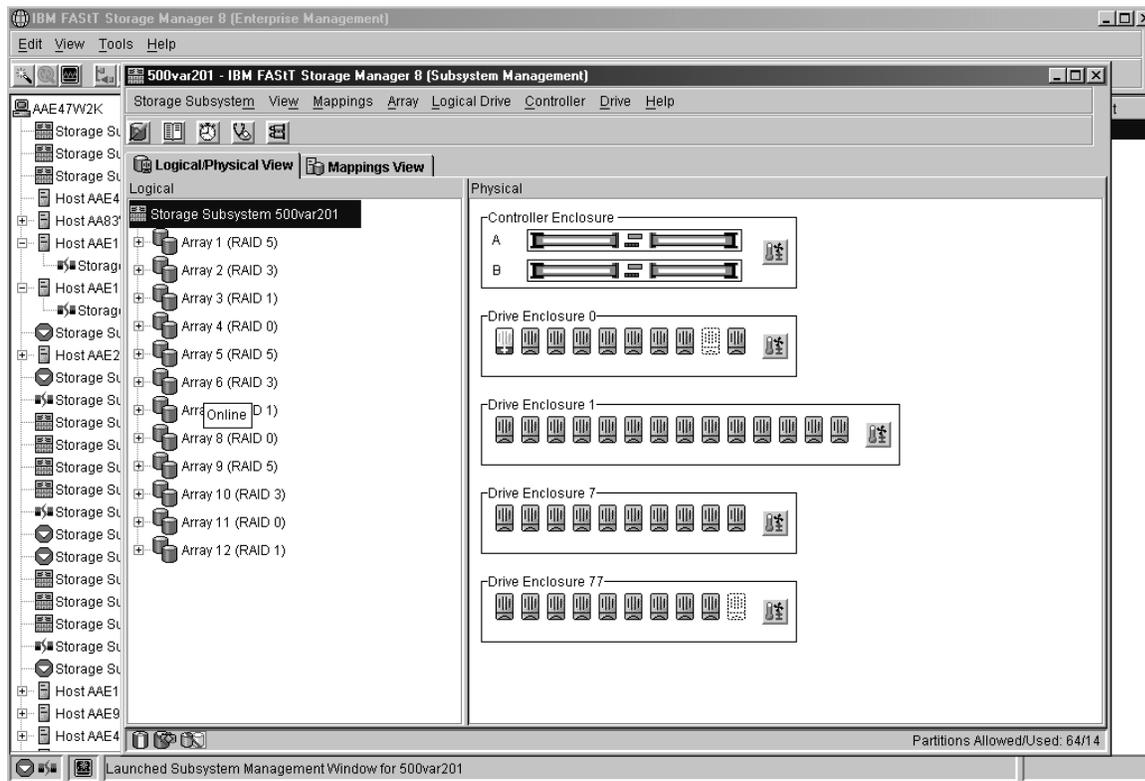

SJ000713

Figure 5. Volume Properties window

Redistributing LUNs in case of failure

If a failure occurs that initiates a controller failover, follow this procedure to
redistribute logical drives to their preferred paths.

1. Repair or replace any faulty components.
2. Complete these steps to mark the failed controller active.
   a. Go to the Subsystem Management window.

b. Right-click **controller A** or **controller B**, and then click **Change —> Mode Active**.

3. Redistribute volumes to their preferred paths by clicking **Subsystem Management —> Storage Subsystem —> Redistribute Logical Drive**.

4. Use the **fget_config** command to verify the active paths, as shown in this example:

```
# fget_config -l dar0
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk1
dac0-hdisk2
dac0-hdisk3
dac1-hdisk4
dac1-hdisk5
dac1-hdisk6
dac1-hdisk7
dac0-hdisk8
```

Creating a direct-attached configuration

Storage Manager 8.21 supports FAStT200, FAStT500, and FAStT700 storage servers in direct-attached AIX (4.3.3 or 5.1) configurations.

To create a direct-attached configuration, you must ensure that:

*   One or two AIX servers can be connected to the FAStT storage server.
    *   FAStT200 can support one AIX server.
    *   FAStT500 can support two AIX servers.
    *   FAStT700 can support two AIX servers.
*   Two server FAStT500 or FAStT700 configurations require four host-side mini-hubs, each with exactly one FC connection from each HBA to a mini-hub.
*   There are two or four HBAs (FC 6227 or FC 6228) per FAStT storage server. Each pair must be configured to one FAStT partition.
*   No external hubs are being used.

To set up a direct-attached configuration, follow these steps:

1. Connect the HBAs to each controller or mini-hub port of the FAStT storage server.
2. Start the system.
3. Configure the devices as usual.

Chapter 3. Installing storage management station software on HP-UX systems

This chapter provides the following specific information for HP-UX operating systems.

- Hardware and firmware requirements
- JRE runtime installation
- SMclient installation
- SMagent installation
- SMutil installation
- Storage partitioning for heterogeneous environments
- Enabling multipath I/O with PV-links

Hardware and firmware requirements

For system-level hardware and software requirements, see Appendix B, "HP-UX system requirements" on page 71.

Table 6 lists the supported versions of hardware to use with Storage Manager 8.21.

Table 6. Supported versions of hardware for HP-UX systems

Product Name	Model	Product release and firmware version
IBM FAStT700 RAID Controller Enclosure Unit	1742-1RU	Firmware 05.21.05.00, NVSRAM CNV1742R821NT008
IBM FAStT700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9319
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Firmware 05.21.05.00, NVSRAM CNV3352R821NT008
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9163, 9165
IBM FAStT200 RAID and Storage Unit, Single Controller	3542-1RU	Firmware 05.20.08.00, NVSRAM CNV3542R821NT017
IBM FAStT200 RAID and Storage Unit, Single Controller	3542-1RU and FC 2101	Firmware 05.20.12.00, NVSRAM CNV3542R821NT018
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Firmware 05.20.12.00, NVSRAM CNV3542R821NT018
IBM HA Emulex LP7000	FC 6227	3.22A1
IBM HA Emulex LP9000	FC 6228	3.82A1
HP FC Adapter	A5158A	B.11.00.06 (driver)
HP FC Adapter 2 Gb	A6795A	B.11.00.01 (driver)
Brocade switch	FC 2109	2.6.0.c/3.0.2.c
McData switch	IBM 2031-16	02.00.00.25
McData switch	IBM 2032-001	02.00.00.25
McData switch	IBM 2032-064	02.00.00.25
INRANGE switch	IBM 2042	3.2.1.J

If the controllers are shipped with a firmware release prior to version 05.20.08.00 and an NVSRAM release prior to version CNV3542R821NT016 (for FAStT200 storage servers) or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV3552R821NT007 (for FAStT500 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV1742R821NT007 (for FAStT700 storage servers), they must be upgraded. Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700

Installing the host software on HP-UX hosts

Use the following procedure to install the host software on an HP-UX storage management station or on an HP-UX host acting as a storage management station. Install the software in the following order:

1. JRE
2. SMclient

The host software for HP-UX consists of the following three packages:

1. FAStT Client (SMclient)
2. FAStT Host Agent (SMagent)
3. FAStT utilities (SMutil)

All three packages are dependent on JRE runtime (JRE), which is a Java compiler for the SMclient, SMagent and SMutil. Because the host software is dependent on Java Runtime Environment, failure to install JRE first results in errors during the software install analysis phase.

Prerequisites

Before installing the client software, ensure that the following conditions are met:

- This is the HP 9000-series server that you identified as the storage management station.
- This machine meets the minimum hardware and software requirements described in Appendix B, "HP-UX system requirements" on page 71.
- Neither the JRE, SMclient, SMagent, or SMutil is installed on the host and you are ready to install it on this machine.

Installing JRE on HP-UX hosts

In the following procedure, the installation CD is mounted at /cdrom. Adjust these instructions as required for your specific installation. No restart is required during the JRE installation process.

1. Insert the HP-UX installation CD in the CD-ROM drive.
2. Start the installation process:

```
# mount -o cdcase -F cdfs /<put device pathname here> /cdrom
# swinstall -s /cdrom/HP-UX/HPruntime_0804500
```

The Software Selection window opens and displays this message:

```
JRE
```

3. Select **JRE** to install the Java Runtime Environment (JRE).
4. From the toolbar, click **Actions —> Install (Analysis)**.
5. Click **OK** when the analysis is finished.
6. In the Confirmation window, click **Yes** to start the installation.
7. When the installation is finished, click **Done** in the Install window and close the application.

   Note: You must exit the **swinstall** command before running the **swverify** command to to verify software installation.

8. Verify that the installation was successful:

```
# swverify -v JRE
```

   If no failure is reported, go to step 9. If a failure is reported, follow the instructions in the /var/adm/sw/swagent.log file.

9. You have completed the JRE installation on this HP-UX storage management station.

Installing the SMclient on HP-UX hosts

In the following procedure, the installation CD is mounted at /cdrom. Adjust these instructions as required for your specific installation. No restart is required during the client installation process.

1. Start the installation process:

```
# mount -o cdcase -F cdfs /<put device pathname here> /cdrom
# swinstall -s /cdrom/HP-UX/SMclient-HP-081G500
```

The Software Selection window opens and displays this message:

```
JRE
SMclient
```

2. Select **SMclient** to install the Storage Manager Client.
3. From the toolbar, click **Actions —> Install (Analysis)**.
4. Click **OK** when the analysis is finished.
5. In the Confirmation window, click **Yes** to start the installation.
6. When the installation is finished, click **Done** in the Install window and close the application.

   Note: You must exit the **swinstall** command before running the **swverify** command to to verify software installation.

7. Verify that the installation was successful:

```
# swverify -v SMclient
```

   If no failure is reported, go to step 8 on page 32. If a failure is reported, follow the instructions in the /var/adm/sw/swagent.log file.

8. You have completed the SMclient software installation on this HP-UX storage management station.

Note: To ensure redundancy in a cluster environment, you must install the client software on at least one additional storage management station or cluster server.

Installing the client software on HP-UX hosts

Use the following procedure to install the client software.

Install the software in the following order:
1. SMagent
2. SMutils

Installing the SMagent on HP-UX hosts

In the following procedure, the installation CD is mounted at /cdrom. Adjust these instructions as required for your specific installation. No restart is required during the SMagent installation process.

1. Start the installation process:

```
# mount -o cdcase -F cdfs / <put device pathname here> /cdrom
# swinstall -s /cdrom/HP-UX/SMagent-HP-08204500
```

The Software Selection window opens and displays this message:

```
JRE
SMagent
```

2. Select **SMagent** to install the Storage Manager Agent.
3. From the toolbar, click **Actions —> Install (Analysis)**.
4. Click **OK** when the analysis is finished.
5. In the Confirmation window, click **Yes** to start the installation.
6. When the installation is finished, click **Done** in the Install window and close the application.

   Note: You must exit the **swinstall** command before running the **swverify** command to to verify software installation.
7. Verify that the installation was successful:

```
# swverify -v SMagent
```

If no failure is reported, go to step 8. If a failure is reported, follow the instructions in the /var/adm/sw/swagent.log file.
8. You have completed the SMagent software installation on this HP-UX storage management station.

Installing the SMutils on HP-UX hosts

In the following procedure, the installation CD is mounted at /cdrom. Adjust these instructions as required for your specific installation. No restart is required during the SMutils installation process.

1. Start the installation process:

```
# mount -o cdcase -F cdfs /<put device pathname here> /cdrom
# swinstall -s /cdrom/HP-UX/SMutil_08204500
```

   The Software Selection window opens and displays this message:

```
JRE
SMutils
```

2. Select **SMutils** to install the Storage Manager utilities.
3. From the toolbar, click **Actions —> Install (Analysis)**.
4. Click **OK** when the analysis is finished.
5. In the Confirmation window, click **Yes** to start the installation.
6. When the installation is finished, click **Done** in the Install window and close the application.

   Note: You must exit the **swinstall** command before running the **swverify** command to to verify software installation.

7. Verify that the installation was successful:

```
# swverify -v SMutils
```

   If no failure is reported, go to step 8. If a failure is reported, follow the instructions in the /var/adm/sw/swutil..log file.
8. You have completed the SMutil software installation on this HP-UX storage management station. Remove the installation CD from the CD-ROM drive.

Configuring storage and partitioning for heterogeneous environments

After you install Storage Manager 8.21 software you need to configure the subsystem. System configuration includes configuring logical drives through the Subsystem Management window of the SMclient application.

For more information about configuring storage see *IBM TotalStorage FAStT700 and Copy Service*, an IBM Redbook publication.

See "Configuring storage subsystems" on page 58 for additional information on storage configuration.

Note: An HP-UX system cannot be in a heterogeneous environment that includes AIX hosts.

Read the following before continuing:

- All FAStT storage servers ship with NVSRAM configured for Windows NT hosts and have an access volume set to LUN 31.
- AVT/ADT is enabled by default; leave it in that state for HP-UX hosts.
- If the controllers are shipped with a firmware release prior to version 05.20.08.00 and an NVSRAM release prior to version NV3542R810N016 (for FAStT200 storage servers), or NV3552R821NT007 (for FAStT500 storage servers), or a firmware release prior to 05.21.01.01 and an NVSRAM release prior to CNV1742R821NT007 (for FAStT700 storage servers), they must be upgraded. Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

www.storage.ibm.com/fast200/supserver.htm
www.storage.ibm.com/fast500/supserver.htm
www.storage.ibm.com/fast700/supserver.htm

> **Note:** The versions of firmware and NVSRAM that ship with this product can be later releases than those described in this document.

Configuring storage

Complete the following procedure to configure Storage Manager 8.21 for an HP-UX system.

1. To set up the storage subsystem for HP-UX, the subsystem must be physically configured, at least initially, for direct management through the Ethernet connections on each controller as well as through the fibre-channel connection. Install the Storage Manager 8.21 software (JRE, SMclient, SMagent, and SMutils) before configuring the subsystem.

2. After the disk subsystem is configured on the network, type the following command to run the SMclient software on either the host server, if it is on the same network as the storage subsystem, or on another machine in the network that has the Storage Manager 8.21 software installed:

```
# SMclient
```

3. Complete the following steps to specify the IP addresses of the controllers.

   a. In the Enterprise Management window, click **Edit —> Add Device**.

   b. In the Add Device window, type the IP address of the first controller in the storage subsystem and click **Add**.

   c. Type the IP address of the second controller and click **Add**, and then click **Close**.

4. In the Subsystem Management window, click **Storage Subsystem —> Profile** and review the summary portion of the output.

   Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

   www.storage.ibm.com/fast200/supserver.htm
   www.storage.ibm.com/fast500/supserver.htm
   www.storage.ibm.com/fast700/supserver.htm

5. Complete the following steps to upgrade the NVSRAM.

   a. Rescan the storage subsystem, and then click **Manage the Device** to return to the Subsystem Management window.

   b. Click **Storage Subsystem —> Download —> NVSRAM**.

   c. In the NVSRAM window, go to the directory where the latest NVSRAM file resides.

   d. Type or select the full pathname of the NVSRAM file, and then click **Update —> OK**.

   e. When the NVSRAM download is finished, turn the power to the storage subsystem off and on in the following sequence.

      1) Turn off both controllers.

      2) Turn off any attached drive enclosures.

      3) Wait at least 30 seconds.

      4) Turn on any attached drive enclosures.

      5) Turn on both controllers.

If required, the host restart process starts after the storage subsystem has been turned off.

6. Complete the following steps to upgrade the firmware.

   a. Click **Storage Subsystem —> Download —> Firmware**.

   b. Go to the directory where the firmware resides.

   c. Type or select the full pathname of the firmware file, and then click **Update —> OK**.

   The firmware is downloaded to the controllers. When the download is finished, the storage subsystem becomes unresponsive and you are returned to the Enterprise Management window.

7. Type the following command to restart the host:

```
# reboot
```

Partitioning storage

Note: Before you partition the storage, identify the worldwide port name for each of your HBAs. Look for the world wide port name for the HBA and write down the number for future reference. Use the following command to find your HBAs listed in the dev directory as `td0`, `td1`, and so on, where *X* in the following example represents the td number.

```
#cd /dev/td
#tdutil /dev/td<X>
```

1. Click the **Mappings View** tab on the Subsystem Management window.

2. In the Mappings window, create a new host group by clicking **Mappings —> Define —> Host Group**.

3. Type the name of the new host group (for example, `HP-UX`). Click **Add**, and then click **Close**.

4. Highlight the new host group and click **Mappings —> Define —> Host**.

5. Define the new host. Type the name of the HP-UX host to which the storage subsystem is attached.

   a. Click **Add**, and then click **Close**.

   b. Highlight the host that you just added and right-click **Define New Host Port**.

      Note: You might be required to restart the host if you cannot see the WWPN for your HBAs. Restarting causes the system to log into the switch.

   c. Select the desired host port for the first HBA, and then change the host type to **HP-UX** and click **Add**.

   d. Choose the host port for the second HBA and click **Add**, and then click **Close**.

Enabling multipath I/O with PV-links

If the HP-UX system is attached with two host bus adapters to the FAStT storage server, you can establish redundant access to storage by using physical volume links (PV-links), a feature of the HP-UX operating system. PV-links achieve access redundancy by using devices with both primary and secondary paths to the same device. Use the following procedure to establish redundant access to storage by using PV-links.

1. Run the **hot_add** command from HP-UX at the shell prompt. This command updates any new devices that are created or added. A dump will is generated. When the **hot_add** utility runs, each new logical drive that iscreated in the Subsystem Management window represents a disk device to the operating system.

```
#hot_add
```

2. Run the **SMdevices** command. The system provides a dump similar to the example in Figure 6 on page 37. Notice that every logical drive or device and access logical unit number (LUN) is listed twice, because you have two paths to each logical drive.

```
#SMdevices
```

```
IBM FAStT Storage Manager Devices, Version 08.20.45.00
Built Wed Mar 20 00:58:59 GMT+00:00 2002
(C) Copyright International Business Machines Corporation, 2002
Licensed Material - Program Property of IBM.  All rights reserved.



/dev/rdsk/c28t0d2 [Storage Subsystem 700var201, Logical Drive hp_1, LUN 1,
Logical Drive WWN <600a0b8000075a54000000803d2c6de6>, Preferred Path
(Controller-A): In Use]

/dev/rdsk/c28t0d3 [Storage Subsystem 700var201, Logical Drive hp_2, LUN 2,
Logical Drive WWN <600a0b8000075a540000006d3d2c6d9f>, Alternate Path
(Controller-A): In Use]

/dev/rdsk/c28t0d4 [Storage Subsystem 700var201, Logical Drive hp_3, LUN 3,
Logical Drive WWN <600a0b8000075a54000000813d2c6df4>, Preferred Path
(Controller-A): In Use]

/dev/rdsk/cc28t0d5 [Storage Subsystem 700var201, Logical Drive hp_4, LUN 4,
Logical Drive WWN <600a0b8000075a540000006e3d2c6da9>, Alternate Path
(Controller-A): In Use]

/dev/rdsk/c28t3d7 [Storage Subsystem 700var201, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b8000075a540000000200000000>]


/dev/rdsk/c30t0d2 [Storage Subsystem 700var201, Logical Drive hp_1, LUN 1,
Logical Drive WWN <600a0b8000075e60000000803d2c6de6>, Preferred Path
(Controller-A): In Use]

/dev/rdsk/c30t0d3 [Storage Subsystem 700var201, Logical Drive hp_2, LUN 2,
Logical Drive WWN <600a0b8000075e600000006d3d2c6d9f>, Alternate Path
(Controller-A): In Use]

/dev/rdsk/c30t0d4 [Storage Subsystem 700var201, Logical Drive hp_3, LUN 3,
Logical Drive WWN <600a0b8000075e60000000813d2c6df4>, Preferred Path
(Controller-A): In Use]

/dev/rdsk/c30t0d5 [Storage Subsystem 700var201, Logical Drive hp_4, LUN 4,
Logical Drive WWN <600a0b8000075e600000006e3d2c6da9>, Aternate Path
(Controller-A): In Use]

/dev/rdsk/c30t3d7 [Storage Subsystem 700var201, Logical Drive Access, LUN 31,
Logical Drive WWN
<600a0b8000075e600000000200000000>]
```

Figure 6. Device names

3. Determine the primary and alternate path for each logical drive. A part of the WWN of each logical drive is unique for each controller in the FAStT storage server. If you examine the WWNs for the access volumes in Figure 6, you will notice that they differ in only five digits, `75a54` and `75e60`. The WWNs for the logical drives also differ in these five digits.

   In the example in Figure 6, the devices are viewed through the adapters `c28` and `c30`. To determine the primary path of a specific logical drive do the following:

   a. Find the WWN for each access LUN. In this case, Access LUN 1 is associated with c28 and has the WWN of 75a54. Access LUN 2 is associated with c30 and has the WNN of 75e60.

    b. Identify the primary path by matching the logical drive WWN to an access LUN WWN. In this case, the WWN for the LUN 10, (75a54) is associated with adapter c28. Therefore, the primary path would be c28t1d2 and the alternate path would be c30t1d2.

    c. For future reference, enter this path information into a primary and secondary path matrix, similar to the one in Table 7.

Table 7. Primary and secondary path matrix information

Logical drive	Primary path	Secondary path
hp_1	c28t1d2	c30t1d2

    d. Repeat these steps for each logical drive.

4. Define the primary paths. The system confirms the creation of the new physical volume.

```
#pvcreate /dev/rdsk/c28t0d0
```

5. Make a directory for volume group. You need the directory to reside in the /dev directory.

```
#cd /dev
#mkdir vg1
```

6. Make a node.

```
#mknod /dev/vg1/group c 64 0x01000
```

7. Create PV-links for the device with the primary path.

```
#vgcreate /dev/vg1 /dev/dsk/C28t1d2
```

8. Add an alternate path for the device.

```
#vgextend vg01 /dev/dsk/C30t1d2
```

9. Repeat steps 4 - 8 for each device.

10. Verify the primary and alternate paths for each device, where *vgname* is the volume group name.

```
#vgdisplay -v <vgname>
```

Identifying the controller ID numbers

Complete this procedure to identify the controller ID numbers:

1. Select a controller in the Physical View of the Subsystem Management window.

2. Click **Controller —> Properties —> Interfaces**. The Controller Properties window opens, as shown in Figure 7 on page 39.
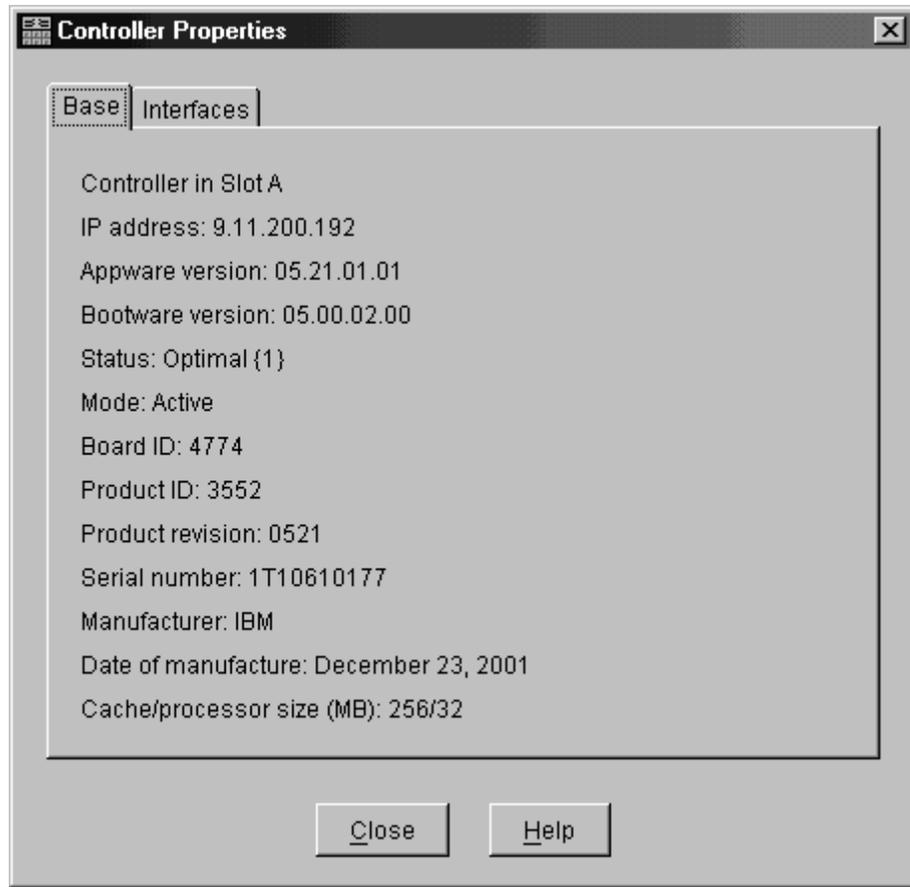
Figure 7. Controller Properties window

3. The first number in the **Current Loop ID** field is the controller ID, which can be one or two digits. For hubs and directly-connected subsystems, the controller ID number and the switch port ID can be found in the same window by converting the middle hexadecimal number to a decimal number.

Identifying device names and bus numbers

Use the **ioscan** command to determine the full device names of data logical drives in the storage subsystem. This information allows you to determine the host adapter external bus numbers and the IDs of the controllers that the host adapters are connected to.

1. View the operating system list of devices:

```
# ioscan -f
```

2. The identification information for every device is displayed. See Figure 8 on page 40.

```
Class           I     H/W Path            Driver   S/W State   H/W Type     Description
=====================================================================================
bc              0                         root     CLAIMED     BUS_NEXUS
bc              1     8                    ccio     CLAIMED     BUS_NEXUS    I/O Adapter
fc              0     8/12                 fcT1     CLAIMED     INTERFACE    HP Fibre Channel
Stora  External bus number of      Controller ID
lan       the host adapter             number     CLAIMED     INTERFACE    HP Fibre Channel
Storage Cntl
fcp             0     8/12.8               fcp  Target address  INTERFACE    FCP Protocol Adapter
ext_bus         2     8/12.8.0.2.0.0       fcpdev   CLAIMED     INTERFACE    FCP Device Interface
target          39    8/12.8.0.2.0.0       tgt      CLAIMED     DEVICE
disk            130   8/12.8.0.2.0.0.0     sdis  LUN number     DEVICE       IBM INF-01-00
disk            35    8/12.8.0.2.0.0.7     sdisk    CLAIMED     DEVICE       IBM Universal Xport
target          39    8/12.8.0.2.0.1       tgt      CLAIMED     DEVICE
disk            131   8/12.8.0.2.0.1.0     sdisk    CLAIMED     DEVICE       IBM INF-01-00
                          Switch port
                          ID number
                                                                                    SJ000701
```

Figure 8. Device identification information

3. Examine the information. LUN numbers range from 0 - 7 and target addresses from 0 - 3.

   Device names have the following format: c<*x*>t<*y*>d<*z*> where:

   <*x*>
   Represents the controller ID number

   <*y*>
   Represents the target address

   <*z*>
   Represents the LUN number

   For example, disk 130 in Figure 8 has a device name of c2t0d0.

4. Determine the entries for the logical drives using the controller ID and LUN numbers. Find the **ext_bus** row above the logical drive listing and note the bus number in the **I** column.

5. For each storage subsystem, use the Mappings window (see Figure 9 on page 41). Click the **Mappings View** tab in Subsystem Management window.

   • Select the storage subsystem node in the Topology View to see a list of all logical drives in the Mappings View.

   • Use the LUN numbers in the Mappings View to associate the logical drives with their device names.

SJ000703
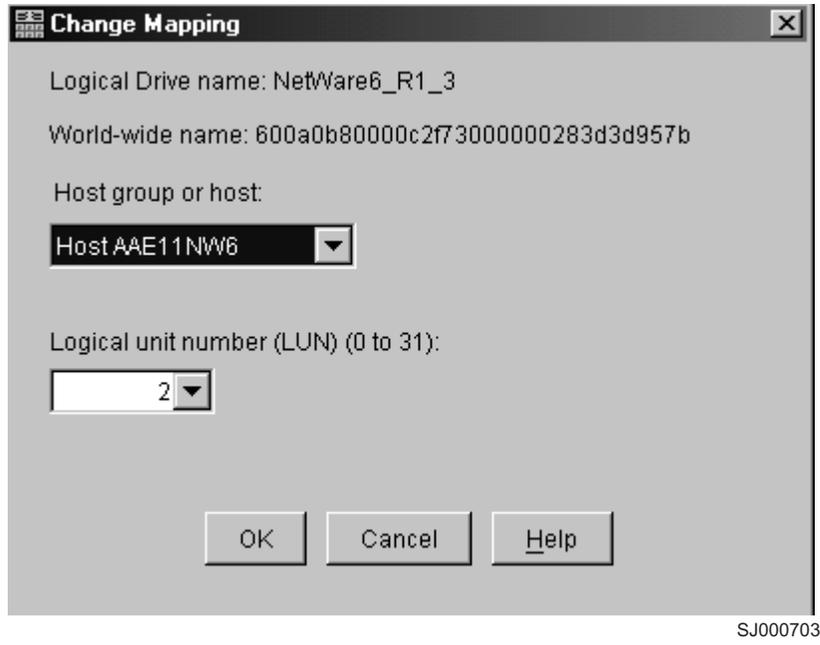
Figure 9. Mappings window

Chapter 4. Installing storage management station software on Solaris systems

This chapter provides the following specific information for Solaris operating systems.

- Hardware and firmware requirements
- Runtime software installation
- Client software installation
- Host software installation
- Configuration information

Hardware and firmware requirements

Table 8 lists the supported versions of hardware to use with Storage Manager 8.21.

Table 8. Supported versions of hardware for Solaris systems

Product Name	Model	Product release and firmware version
IBM FAStT700 RAID Controller Enclosure Unit	1742-1RU	Firmware 05.21.05.00, NVSRAM CNV1742R821NT008
IBM FAStT700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9319
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Firmware 05.21.05.00, NVSRAM CNV3352R821NT008
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9163, 9165
IBM FAStT200 RAID and Storage Unit, Single Controller	3542-1RU	Firmware 05.20.08.00, NVSRAM CNV3542R821NT017
IBM FAStT200 RAID and Storage Unit, Single Controller	3542-1RU and FC 2101	Firmware 05.20.12.00, NVSRAM CNV3542R821NT018
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Firmware 05.20.12.00, NVSRAM CNV3542R821NT018
IBM HA Emulex LP7000	FC 6227	3.22A1
IBM HA Emulex LP9000	FC 6228	3.82A1
JNI FC Adapter	FC64-1063	fcaw2.5.18.pkg
JNI FC Adapter	FCE-1063 S-bus	JNIC.pkg, version 4.1.1.1
JNI FC Adapter	FCE-1063 S-bus 64 bit	JNIC.pkg, version 4.1.1.1
JNI FC Adapter	FCE-1473 S-bus 64 bit	JNIC146x.pkg, version 5.1.2
JNI FC Adapter	FCE-6410-N 33 Mhz	JNIC.pkg, version 4.1.1.1
JNI FC Adapter	FCE-6460-PCI 64 bit	JNIC146x.pkg, version 5.1.2
JNI FC Adapter	FCE2-1063 S-bus 64 bit	JNIC.pkg, version 4.1.1.1
JNI FC Adapter	FCE2-6412 66 Mhz 64 bit	JNIC.pkg, version 4.1.1.1

Table 8. Supported versions of hardware for Solaris systems (continued)

Product Name	Model	Product release and firmware version
JNI FC Adapter	FCI-1063	fca-pci.2.5.18.pkg
Brocade switch	FC 2109	2.6.0.c, 3.0.2.c
McData switch	IBM 2031-16	02.00.00.25
McData switch	IBM 2032-001	02.00.00.25
McData switch	IBM 2032-064	02.00.00.25
INRANGE switch	IBM 2042	3.2.1.J

Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700

For system-level hardware and software requirements, see Appendix C, "Solaris system requirements" on page 73.

Installing the runtime software on Solaris hosts

Use the following procedure to install the runtime software (SMruntime) on a Solaris storage management station.

Prerequisites

Before installing the runtime software, ensure that the following conditions are met:
- This is the SUN SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 2.6, 2.7, or 8 and meets the minimum hardware and software requirements described in Appendix C, "Solaris system requirements" on page 73.
- SMruntime is not installed and you are ready to install it on this machine.

Procedure

In the following procedure, the installation CD is mounted at /cdrom/sm820. Adjust these instructions as required for your specific installation location. No restart is required during the client installation process.

1. Insert the Solaris installation CD in the CD-ROM drive.
2. Start the installation process:

```
pkgadd -d /cdrom/SM820/SUN_SParc/SM8runtime-Sparc<filename>.pkg
```

Information about packages that can be installed in the specified directory is displayed on the command line. The following is an example of what you might see displayed.

```
The following packages are available:
1 SMruntime
IBM FAStT Storage Manager 8 Runtime (sparc)
<version number>
Select package(s) you wish  to process (or 'all' to process all
Packages).  (default:all)   [?,??,q]:
```

3. Type the value of the package that you are installing and press Enter. The installation process begins. The following prompt is displayed.

```
This package contains scripts which will be executed with super-user
Permission during the process of installing this package.
Do you want to continue with the installation of <SMruntime>
[y, n, ?]
```

4. Type *y* and press Enter. The installation process continues. When the SMruntime software has been successfully installed, the following message is displayed:

```
Installation of <SMruntime> was successful.
```

Installing the client software on Solaris hosts

Use the following procedure to install and integrate the client software on a Solaris storage management station.

Note: SMclient is supported on Solaris 2.6, 2.7, and 8.

Installing SMclient on Solaris hosts

Before installing the client software, ensure that the following conditions are met:

- This is the Sun SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 2.6, 2.7 or 8 and meets the other minimum hardware and software requirements described in Appendix C, "Solaris system requirements" on page 73.
- SM runtime is installed.
- SMclient is not installed and you are ready to install it on this machine.

Procedure

In the following procedure, the installation CD is mounted at /cdrom/sm821. Adjust these instructions as required for your specific installation location. No restart is required during the client installation process.

1. Insert the Solaris installation CD in the CD-ROM drive.
2. Start the installation process:

```
# pkgadd -d /cdrom/SM821/SUN_Sparc/SM8client-Sparc<filename>.pkg
```

Information about packages that can be installed in the specified directory is displayed on the command line. The following is an example of what you might see displayed.

```
The following packages are available:

1 SMclient                      IBM FAStT Storage Manager 8 Client
                                 (sparc) <version number>

Select package(s) you wish to process (or 'all' to process all
packages). (default:all)  [?,??,q]:
```

3. Type the value of the package that you are installing and press Enter.

   The installation process begins. The following prompt is displayed:

```
This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <SMclient>

[y, n, ?]
```

4. Type y and press Enter.

   The installation process continues. When the SMclient software has been successfully installed, the following message is displayed:

```
Installation of <SMclient> was successful.
```

5. Type the following command to verify that the installation was successful:

```
# pkgchk SMclient
```

   The installation was successful if no output is displayed on the screen.

6. You are finished with the client software installation on this Solaris storage management station. Remove the installation CD from the CD-ROM drive.

7. Choose one of the following options:

   Storage management software is required on additional storage management stations
   Repeat step 1 on page 45 through step 6, using the corresponding installation profile for each storage management station as a guide.

   Storage management software installation is completed on all storage management stations
   You are now ready to install the host software. Go to "Installing host software on Solaris hosts".

   Note: To ensure redundancy in a cluster environment, you must install the client software on at least one additional storage management station or cluster server.

Installing host software on Solaris hosts

After you install the Storage Manager 8.21 client software and configure your storage subsystems, use these instructions to install the appropriate Storage Manager 8.21 host software. See Table 1 on page 3 for information about the software packages and the installation sequence required by each host type.

All Solaris hosts in your storage subsystem must have the following host software packages installed on them:

- RDAC

- SMagent
- SMutil

This section contains the procedure for installing the software.

Prerequisites

Before installing the host software, ensure that the following conditions are met:
- This is the Sun SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 2.6, 7, or 8, and meets the minimum hardware and software requirements described in Appendix C, "Solaris system requirements" on page 73.
- The RDAC, SMagent, and SMutil software packages are not installed and you are ready to install them on this machine.

Note: If you plan to run Storage Manager 8.21 in an environment that includes storage arrays that you are not upgrading to 8.21, you must upgrade your existing storage manager to at least version 6.22. The Storage Manager 8.21 RDAC package replaces the version 6.22 RDAC files and is shared by version 6.22 and version 8.21 of the storage management software.

See the Storage Manager 6.x installation instructions for information about uninstalling the software.

Installing RDAC on Solaris hosts

In the following procedure, the installation CD is mounted at /cdrom/SM821. Adjust these instructions as required for your specific installation location.

1. Insert the Solaris installation CD in the CD-ROM drive.
2. Type the following command to start installing the RDAC package:

```
# pkgadd -d /cdrom/SM821/SUN_Sparc/SM8rdac-Sparc-<filename>.pkg
```

The installation process begins.

Information about packages that can be installed in the specified directory is displayed on the command line. The following is an example of what you might see displayed.

```
The following packages are available:

1 RDAC                           Redundant Disk Array Controller
                                 (sparc) <version number>

Select package(s) you wish to process (or 'all' to process all
packages). (default:all)  [?,??,q]:
```

3. Type the value of the package you are installing and press Enter.

   The installation process begins.
4. The software automatically checks for package conflicts. If any conflicts are detected, a message is displayed indicating that some files are already installed and are in use by another package.

The following prompt is displayed:

```
Do you want to install these conflicting files [y, n, ?]
```

Type y and press Enter.

5. The following prompt is displayed:

```
This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <RDAC>

[y, n, ?]
```

Type y and press Enter.

The installation process continues.

6. When the RDAC package has been successfully installed, the following message is displayed:

```
Installation of <RDAC> was successful.
```

Note: Make sure that the variables in the configuration files for your JNI adapter cards have been set to the correct values. For information about these configuration files and their variables, see "Variable settings for JNI adapter cards" on page 51.

Installing SMagent on Solaris hosts

1. Type the following command to start installing the SMagent package.

```
pkgadd -d /cdrom/SM821/SUN_SParc/SM8agent-Sparc-<filename>.pkg
```

The installation process begins.

Information about packages that can be installed in the specified directory is displayed on the command line. The following is an example of what you might see displayed.

```
The following packages are available:

1 SMagent                   IBM FAStT Storage Manager 8 Agent
                            (sparc) <version number>

Select package(s) you wish to process (or 'all' to process all
packages). (default:all)  [?,??,q]:
```

2. Type the value of the package that you are installing and press Enter.

The installation process begins. The following prompt is displayed.

```
This package contains scripts which will be executed with super-user
Permission during the process of installing this package.
Do you want to continue with the installation of <SMruntime>
[y n, ?]
```

3. Type y and press Enter.

The installation process continues. When the SMagent software has been successfully installed, the following message is displayed:

```
Installation of <SMagent> was successful.
```

4. Move to the root directory:

```
# cd /
```

5. Stop the system:

```
# halt
```

6. Restart the system with either this command:

```
# boot -r
```

or this command:

```
# boot <bootdevices> -r
```

7. Remove the installation CD from the CD-ROM drive.
8.

You have finished installing the Storage Manager 8.21 software on this Solaris host.

Performing the initial configuration of storage subsystems on Solaris hosts

Read the following before continuing:

- All FAStT storage servers ship with NVSRAM configured for Windows NT hosts and have an access volume set to LUN 31. This should not present any problems for Solaris in recognizing the storage subsystem.
- If the controllers are shipped with a firmware release prior to version 05.20.08.00 and an NVSRAM release prior to version CNV3542R821NT016 (for FAStT200 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV3552R821NT007 (for FAStT500 storage servers), or a firmware release prior to version 05.21.01.01 and an NVSRAM release prior to version CNV1742R821NT007 (for FAStT700 storage servers), they must be upgraded.

  Read the README file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

  www.storage.ibm.com/fast200/supserver.htm
  www.storage.ibm.com/fast500/supserver.htm
  www.storage.ibm.com/fast700/supserver.htm

  Note: Versions of firmware and NVSRAM that ship with this product can be later releases than those described in this document.

If you do not have the latest versions, download them and continue with step 1 on page 50. If the version numbers are current, then go to step 7 on page 50.

Complete the following procedure to configure Storage Manager 8.21 for a Solaris system:

1. Upgrade the NVSRAM and then upgrade the firmware. In the Subsystem Management window, click **Storage Subsystem —> Download —> NVSRAM**.

2. In the NVSRAM window, go to the directory where the latest NVSRAM file resides.

3. Type or select the full pathname of the NVSRAM file, and then click **Update —> OK**.

4. Click **Storage Subsystem —> Download —> Firmware**.

5. Type or select the full pathname of the firmware file, and then click **Update —> OK**.

   The firmware is downloaded to the controllers. When the download is finished, the firmware becomes unresponsive and you are returned to the Enterprise Management window.

6. Turn the power to the storage subsystem off and on in the following sequence.

   a. Turn off both controllers.

   b. Turn off any attached drive enclosures.

   c. Wait at least 30 seconds.

   d. Turn on any attached drive enclosures.

   e. Turn on both controllers.

7. Rescan the storage subsystem, and then click **Manage the Device** to return to the Subsystem Management window.

8. In the Subsystem Management window, click the **Mappings View** tab.

9. In the Mappings window, create a new host group by clicking **Mappings —> Define —> Host Group**.

10. Type the name of the new host group (for example, `Solaris`). Click **Add**, and then click **Close**.

11. Highlight the new host group and click **Mappings —> Define —> Host**.

12. Define the new host. Type the name of the Solaris host to which the storage subsystem is attached.

    a. Click **Add**, and then click **Close**.

    b. Right-click the host that you just added and select **Define New Host Port**.

    c. Select the desired host port for the first HBA, and then change the host type to **Solaris** and click **Add**.

    d. Choose the host port for the second HBA and click **Add**, and then click **Close**.

    e. To view the added LUNs, click the host group, and then click **Mappings —> Define —> Define Additional Mapping**.

    f. Type the following command to run the hot_add utility:

    ```
    # /etc/raid/bin/hot_add
    ```

       For information about the output of the utility, see "Default partitioning for Solaris devices" on page 51.

    g. Click **File —> Exit**.

    h. Obtain the device name for each volume (or LUN):

    ```
    # format
    ```

The SMclient is now able to connect to the storage subsystem.

Uninstalling Storage Manager 8.21 software

When you uninstall SMruntime, RDAC, SMclient, SMagent, or SMutil software, the directory /var/opt/SM8 might not be removed. Delete this directory to ensure a clean installation of the new software.

Default partitioning for Solaris devices

After you install the Storage Manager 8.21 software and configure the subsystems and logical drives through the Subsystem Management window of the SMclient application, run the /etc/raid/bin/hot_add utility from a shell prompt. (For additional information about the hot_add utility see the applicable man page.)

When the hot_add utility is run, each logical drive that is created in the Subsystem Manager presents a disk device to the operating system. The devices have the following naming convention:

```
/dev/[r]dsk/c<X>t<X>d<X>s<X>
        |     |    |    |    |
        |     |    |    |    s=slice
        |     |    |    d=drive
        |     |    t=target
        |     c=controller
        disk subdirectory
```

where X represents a number. The number following the slice (s) represents the partition on the disk and can be from 0 - 7.

Only mount filesystems on slices 3, 4, 5, 6, or 7 on the devices that are derived from the logical drives on the storage subsystem. If raw partitions or slices 0 - 2 are used, disk information can be overwritten, resulting in an inability to mount the disk. Leave the default root and swap partitions as slices 0 and 1.

Configuring JNI host bus adapter cards

The RDAC software includes the following shell scripts that you can use to configure JNI cards for use with connections through fibre-channel switches:

- For FCI-1063 cards, /etc/raid/bin/genfcaconf
- For all jnic.conf and jnic.146x.conf cards, /etc/raid/bin/genjniconf
- For FC64-1063 and all other cards, /etc/raid/bin/genscsiconf, which calls other scripts as appropriate

For additional information, see the man pages for each script.

Variable settings for JNI adapter cards

Table 9 lists the configuration file for each JNI adapter driver package.

Table 9. JNI adapter configuration files

JNI adapter driver package	Configuration file
fcaw2.5.18.pkg	/kernel/drv/fcaw.conf
fca-pci.2.5.18.pkg	/kernel/drv/fca-pci.conf

Table 9. JNI adapter configuration files (continued)

JNI adapter driver package	Configuration file
JNIC.pkg, version 4.1.1.1	/kernel/drv/jnic.conf
JNIC146x.pkg, version 5.1.2	/kernel/drv/jnic146x.conf

Ensure that the configuration file for your adapter meets the following criteria:

- The *FCLoopEnabled* and *FcFabricEnabled* variables must be set appropriately for your configuration; see the configuration file for details.
- If the *fca_nport* variable is in the file, set it to 1.
- If the *ip-disable* variable is in the file, set it to 1.
- In fca*.conf files, set *BusyRetryDelay* to 5000.
- In fca*.conf files, set *scsi_probe_delay* to 5000.
- In fca*.conf files, set *link_recovery_delay* to 1000.
- The lines that contain these variables must be uncommented and set to the specified values:
  - *failover* or *FailoverDelay*, set to 30
  - *JniCreationDelay*, set to 10
  - *def_wwnn_binding*, remove the dollar-sign character ($) from the line
  - *def_wwpn_binding*
  - *def_hba_binding*

    Note:

    1. *def_hba_binding* is for fca.pci.conf and fcaw.conf
    2. When the RDAC software is installed, the *def_hba_binding* variable in the fca-pci.conf and fcaw.conf files is set to nonjni. This is the correct behavior.
    3. fcaw.conf should be changed back to original value -*def_hba_binding*="fcaw*".

In all configuration files, set the following variables to the specified values:

- *target0_hba*, set to jnic0 in the jnic.conf file, or to jnic146x0 in the jnic146x.conf file.
- *target_wwpn*, set to the wwpn of the first controller (see the /var/adm/messages file).
- *target1_hba*, set to jnic1 in the jnic.conf file, or to jnic146x1 in the jnic146x.conf file, fca-pci0 in the fca-pci.conf file and fcaw0 in the fcaw.conf file, fca-pci1 in the fca-pci.conf file and fcaw1 in the fcaw.conf file.
- *target1_wwpn*, set to the wwpn of the second controller (see the /var/adm/messages file).

If you made any changes to the configuration file, type the following command to halt the host system:

```
# halt
```

When the OK prompt displays, type the following command to restart the host system:

```
# boot -r
```

Importance of setting the host type on Solaris

After you install the Storage Manager 8.21 software, you must set the host type to Solaris for the operating system to work properly with the storage subsystem. This should have been done automatically, but you should verify that it occurred.

Click on the **Mappings View** tab and select **Mappings —> Show all Host Port Information**.

In the Mappings window:
1. Create a new host group.
2. Add the host to the newly created host group.
3. Add each port to the host and ensure that the host type is set to Solaris (Sparc).

Chapter 5. Completing the software installation

This chapter describes how to complete the installation of Storage Manager 8.21 by performing these tasks:

- Failover protection
- Changing the host type
- Understanding the implications of using heterogeneous hosts in your storage subsystem
- Configuring the storage subsystem
- Starting the Subsystem Management window
- Setting the IBM SAN Switch 2109

Failover protection

Failover protection is available by using multipath drivers with or without AVT/ADT with Storage Manager 7.10. A multipath driver is an I/O path failover driver that is installed on host computers that access the storage subsystem. AVT/ADT is a built-in feature of the controller firmware that allows logical-drive-level failover rather than controller-level failover. The following failover settings are available:

- Multipath driver software on one or more hosts and AVT/ADT enabled on the storage subsystem
- Multipath driver software on one or more hosts and AVT/ADT disabled on the storage subsystem

  Note: This is the only valid AIX configuration.

- No multipath driver software on one or more hosts and AVT/ADT enabled on the storage subsystem (no failover)

Before using the storage management software for the first time, be aware of the failover options that are available for each operating system, as shown in Table 10.

Table 10. Failover options by operating system

Operating system	Multipath driver name	Supported failover options	
		Multipath driver with AVT/ADT enabled	**Multipath driver with AVT/ADT disabled**
AIX	fcp array		X
HP-UX	HP Logical Volume Manager (LVM)	X	
Solaris	FAStT Storage Manager RDAC	X (see Note)	X
Note: Disable AVT/ADT if a large number of logical drives in a Solaris-only environment causes long startup times.			

In this scenario, a pair of active controllers is located in a storage subsystem. When you create a logical drive, you assign a controller to own the logical drive (called the preferred controller) and to control the I/O between the logical drive and the application host along the I/O path. The preferred controller normally receives the I/O requests to the logical drive. However, if a problem along the data path (like a component failure) causes an I/O to fail, the multipath driver issues the I/O to the alternate controller.

When AVT/ADT is enabled and used with a host multipath driver, it ensures that an I/O data path is available for the storage subsystem logical drives. The AVT/ADT feature changes the ownership of the logical drive that is receiving the I/O to the alternate controller. After the I/O data path problem is corrected, the preferred controller reestablishes ownership of the logical drive as soon as the multipath driver detects that the path is working again.

Attention: AVT/ADT is enabled on the storage subsystem as a global setting, so its setting applies for every host that connects to the storage subsystem. Keep this in mind if you are using heterogeneous hosts with your storage subsystem. Hosts in a heterogeneous environment should not share logical drives. For more information about heterogeneous hosts, see Chapter 6, "Setting up the storage subsystem in a heterogeneous environment" on page 65.

No multipath driver with AVT/ADT enabled

In this scenario, storage subsystems have no failover protection. A pair of active controllers can still be located in a storage subsystem and each logical drive on the storage subsystem can be assigned a preferred controller. However, logical drives do not move to the alternate controller because there is no multipath driver installed. When a component in the I/O path fails, such as a cable or the controller itself, I/O cannot get through to the storage subsystem. The component failure must be corrected before I/O can resume. (Switching logical drives to the alternate controller in the pair must be done manually.)

Connect hosts that are using operating systems without failover capability to the storage subsystem so that each host adapter has only one path to the controller. With no multipath driver and AVT/ADT enabled, you do not have multiple host bus adapters in the host server. AVT/ADT in this instance is something that can be enabled or disabled on the storage subsystem; for example, you would disable AVT/ADT if you only had one controller. See the SM7client online help for additional information.

Disabling AVT/ADT

On a new storage subsystem, or any time that the NVSRAM has been updated, AVT/ADT is enabled by default. Use the following procedure to disable AVT/ADT:

1. Insert the installation CD in the storage management station CD-ROM drive.
2. Start the storage management software.
3. In the Enterprise Management window, highlight the storage subsystem on which you want to disable AVT/ADT.
4. Click **Tools —> Execute Script**.

   The Script Editor window opens.
5. Click **File —> Load Script**.

   The Load Script window opens.
6. Select avtdisable.scr from the /cdrom/sm710/Scripts directory on the CD.
7. Click **Open** to load the script.
8. Click **Tools —> Execute Only**.

   The script runs. A `Script Execution Complete` message is displayed when the script is finished.
9. Turn off the power to the controllers, then turn the power back on.
10. Go to "Changing the host type" on page 57.

Attention: If you have defined a host array in which more than one host is concurrently accessing the same array on a storage subsystem, consider disabling AVT/ADT. If AVT/ADT remains enabled in this situation, the storage subsystem still functions, but if path failover occurs there are undesirable effects on performance.

For example, say two hosts share access to a logical drive. If one of the hosts experiences a path failure, the controller ownership of the logical drive switches from the original controller owner to the second controller in the pair. Meanwhile, the second host continues to send I/O to the original controller owner, unaware of the path problem that the other host is experiencing. If the I/O data path problem (experienced by the first host) is not related to the controller itself and the I/O reaches the storage subsystem, controller ownership transfers back to the original controller owner. As long as there is a path problem, the hosts continue to transfer controller ownership back and forth with every I/O sent, significantly degrading storage subsystem performance.

Enabling AVT/ADT

If AVT/ADT has been disabled on a storage subsystem, use the following procedure to enable it.

Attention: Refer to Table 10 on page 55 before enabling AVT/ADT on a storage subsystem.

1. Insert the installation CD in the storage management station CD-ROM drive.
2. Start the storage management software.
3. In the Enterprise Management window, highlight the storage subsystem on which you want to enable AVT/ADT.
4. Click **Tools —> Execute Script**.

    The Script Editor window opens.
5. Click **File —> Load Script**.

    The Load Script window opens.
6. Select avtenable.scr from the /cdrom/sm710/Scripts directory on the CD.
7. Click **Open** to load the script.
8. Click **Tools —> Execute Only**.

    The script runs. A `Script Execution Complete` message is displayed when the script is finished.
9. Turn off the power to the controllers, then turn the power back on.
10. Go to "Changing the host type".

Setting the Residual Reservation parameter

On a new storage system, or any time that the NVSRAM has been updated, the Residual Reservation parameter in NVSRAM needs to be set or reset. This is required in high-availability configurations, and when you use High Availability Cluster Multiprocessing (HACMP), Parallel System Support Programs (PSSP), or General Parallel File System (GPFS) software on pSeries or AIX servers.

Changing the host type

The host type defines how the controllers in the storage subsystem work with the operating systems on the hosts that are connected to it.

When using the storage management software, you must set the correct host type by clicking **Mappings —> Change —> Host Type** in the Subsystem Management window.

If partitioning is enabled, click **Mappings —> Change —> Host Type**. Select the desired host port in the left window and click **Configure —> Topology —> Change Host Type**.

For more information, see the topic on changing the host type of a storage subsystem in the Subsystem Management window online help.

If the Storage Partitioning feature is enabled, you must change the host type that is associated with each host port in the Mappings window. (For more information, see the topic on changing the host type of an individual host port in the Subsystem Management window online help.)

When you have selected the host type, continue with "Configuring storage subsystems".

Heterogeneous hosts overview

The heterogeneous hosts feature enables hosts that are running different operating systems to access a single storage subsystem. In previous releases of Storage Manager, only hosts running the same operating system could access a single storage subsystem. Storage Manager 8.21 supports up to 64 storage partitions, which enables a multiple host-type subsystem to share storage capacity, consolidate storage, and reduce storage management costs.

Host computers can be running completely different operating systems (for example, AIX and Solaris) or variants of the same operating system (for example, Solaris running in a cluster environment). When a host type is specified in the Define New Host Port window, the heterogeneous hosts feature enables the controllers in the storage subsystem to tailor their behavior (such as LUN reporting and error conditions) to the needs of the operating system or variant of the host that is sending the information.

Note: The heterogeneous hosts feature is available only with storage partitioning enabled.

In a heterogeneous environment, you must set each host type to the appropriate operating system during host-port definition so that the firmware on each controller can respond correctly for the operating system for that host. Before you begin setting up your heterogeneous hosts configuration (described in Chapter 6, "Setting up the storage subsystem in a heterogeneous environment" on page 65), see the *IBM FAStT Storage Manager Concepts Guide* available on your installation CD or at the following Web site:

www.ibm.com/pc/support/

Configuring storage subsystems

Use the following procedures to start the client software from either the storage management station or from a host that is acting as a storage management station (a host with the client software installed). Use the client software to configure each attached storage subsystem.

Attention: For cluster configurations, complete all applicable configuration procedures for each storage subsystem *before* installing the storage management software on a second host or cluster server.

Starting the Enterprise Management window

Use this procedure to start the Enterprise Management window from the storage management station where you have installed the client software.

1. Start the client software:
   - If you are using a Windows workstation, click **Start —> Programs —> FAStT Storage Manager 8 Client**.
   - If you are using a UNIX-based system, type `SMclient` at a shell prompt.

   A splash screen is displayed while the client software starts. When the client software has been loaded, the Enterprise Management window and the Initial Automatic Discovery window opens (see Figure 10).

   Note: The Enterprise Management window can take several minutes to open. No wait cursor, such as an hourglass, is displayed.
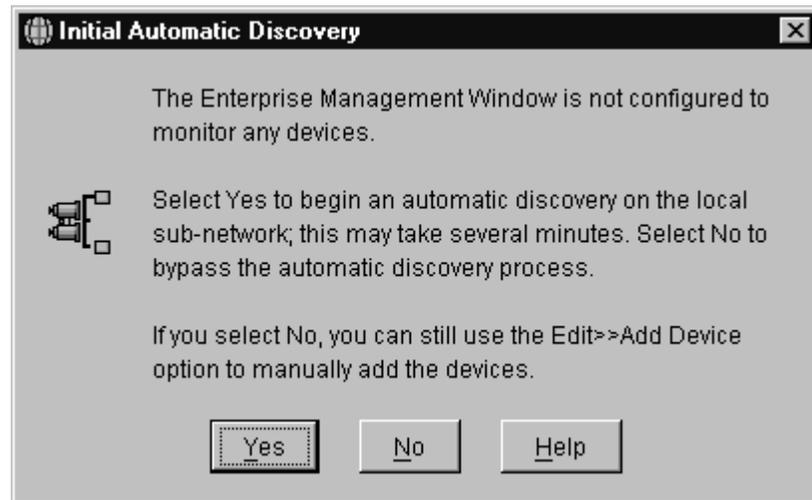
SJ000704

Figure 10. Initial Automatic Discovery window

2. If you are running the client software from AIX, change to the /SMclient directory and then type `SMclient`.
3. If you are running the client software from Solaris 8, type `SMclient`.
4. Click **Yes** to begin an initial automatic discovery of all attached hosts and storage subsystems that are attached.

   The software sends a broadcast message across the local subnetwork that is connected to the storage management station. It discovers host-agent-managed storage subsystems if the respective hosts respond to the broadcast. The software discovers directly managed storage subsystems if the controllers in the attached storage subsystems respond to the broadcast message.

   It can take up to one minute for the Enterprise Management window to refresh after an initial automatic discovery. If you need to stop the automatic discovery operation for any reason, close the Enterprise Management window.

When the initial automatic discovery is finished, all attached hosts and attached storage subsystems are displayed in the Enterprise Management window (see Figure 11).
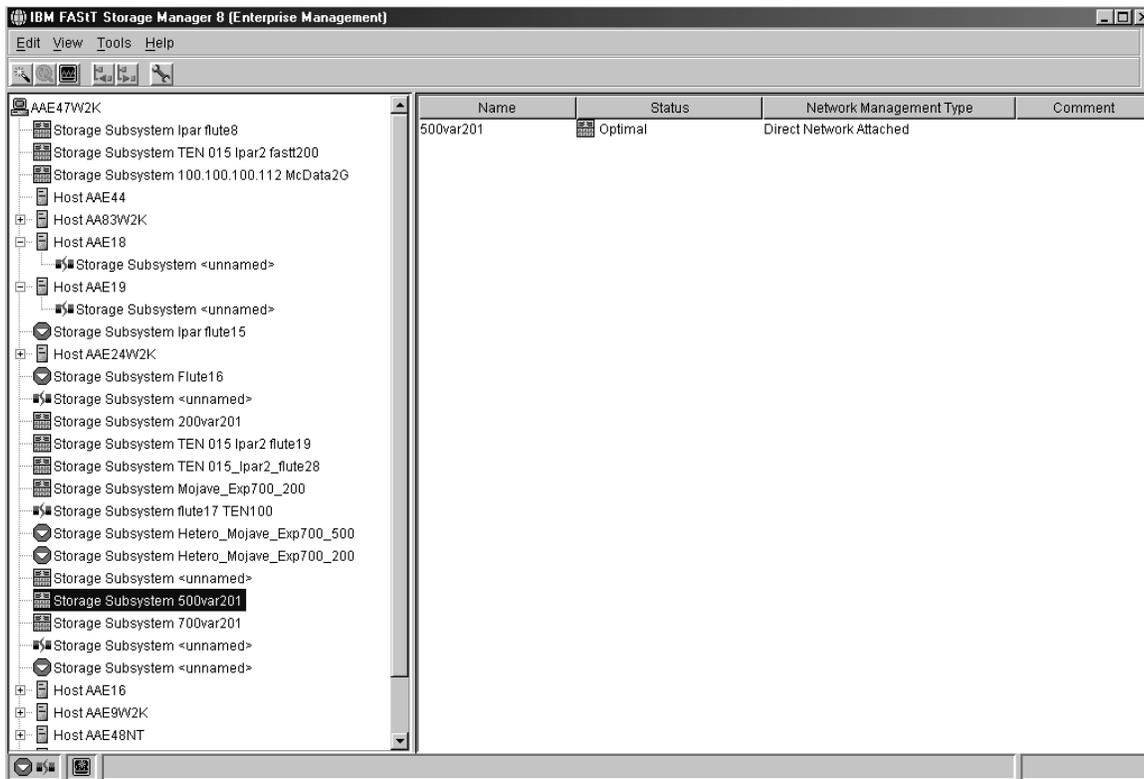
SJ000705

Figure 11. Enterprise Management window

5. Make sure that all the attached hosts and storage subsystems are displayed as expected. If not, do the following:

   a. Check the hardware and connections for possible problems (see the hardware documentation for specific procedures).

   b. See the Enterprise Management window help topic on discovering storage subsystems and take the appropriate action that is indicated.

   c. Make sure that the device is on the local subnetwork. If it is not, click **Edit —> Add Device** to add it. See the Enterprise Management window help topic on adding devices.

   d. If a storage subsystem is duplicated in the device tree after an automatic discovery, remove the duplicate storage subsystem icon from the device tree by clicking **Edit —> Remove Device** in the Enterprise Management window.

6. Make sure that the status of each storage subsystem is **Optimal**.

   • If any device shows a status of **Unresponsive**, remove the device from the management domain, and then add it again. See the Enterprise Management window help topic on removing and adding devices.

   • If the device still shows an unresponsive status, contact your customer service representative.

7. Go to "Configuring alert notifications" on page 61.

Configuring alert notifications

After you have added devices to the management domain, it is recommended that you set up alert notification options to report critical events on the storage subsystems. The following alert notification options are available:

- Alert notifications are sent to a designated network management station (NMS) using simple network management protocol (SNMP) traps.
- Alert notifications are sent to a designated e-mail address. See the Enterprise Management window help for specific procedures. To send e-mail to IBM, contact your customer service representative.
- Alert notifications are sent to a designated alphanumeric pager when third-party software is used to convert e-mail messages. See the Enterprise Management window help for specific procedures.

For more information about notification options, see *IBM FAStT Storage Manager Concepts Guide*, available on your installation CD or at the following Web site:

www.ibm.com/pc/support/

Use the following procedure to set up alert notifications using SNMP traps. You need to set up the designated management station only once.

1. Make sure that the installation CD is inserted in the CD-ROM drive on your designated NMS.
2. From the installation CD, copy the SM8_21.MIB file from the SM8mib directory to the NMS.
3. Follow the steps required by your NMS to compile the MIB. For details, contact your network administrator or see the documentation for the storage management product you are using.
4. Go to "Starting the Subsystem Management window".

Starting the Subsystem Management window

In the Enterprise Management window, select a storage subsystem in one of these ways:

- Click **Tools —> Manage Device**.
- Right-click the **storage subsystem** and click **Manage Device**.
- Double-click the **storage subsystem**.

The Subsystem Management window for the selected storage subsystem opens (see Figure 12 on page 62).

Note: You can manage one storage subsystem per Subsystem Management window. Open multiple windows to manage other storage subsystems.
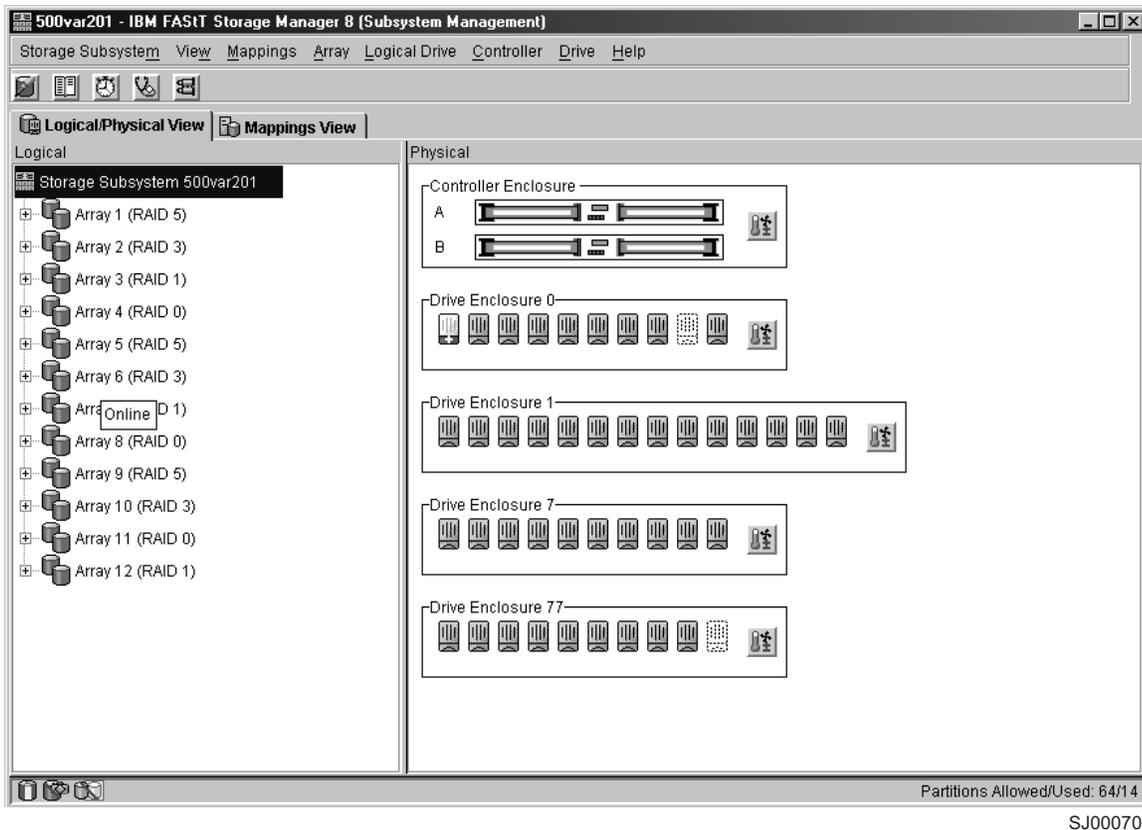
Figure 12. Subsystem Management window

Renaming storage subsystems

When you start the storage management software for the first time, the storage subsystems are unnamed. Rename each storage subsystem from <unnamed> to a name that is appropriate for your network. In the Subsystem Management window, click **Storage Subsystem —> Rename**.

Performing optional storage subsystem management tasks

You can also perform other storage subsystem management tasks at this time. To perform the following tasks and other storage subsystem management tasks, see the appropriate topics in the Subsystem Management window help.

- Locate a storage subsystem
- View a storage subsystem profile
- Configure a storage subsystem password
- Create and manage arrays and array groups
- Use the performance monitor
- Create storage partitions (if applicable)

If you are installing the storage management software in a cluster environment, repeat all applicable client software installation procedures for a second storage management station and for all applicable host software on a second host. See Chapter 2, "Installing storage management station software on AIX systems" on page 11, Chapter 3, "Installing storage management station software on HP-UX

systems" on page 29, or Chapter 4, "Installing storage management station software on Solaris systems" on page 43 for installation instructions.

Setting IBM SAN switches

This section contains information about setting IBM SAN switches to work correctly with the Storage Manager 8.21 in any UNIX or AIX environment.

Keep the following points in mind when configuring the settings on the switch:

- Ensure that the switch is set with an IP address. The IP address can be set using the Ethernet port of the switch or the front panel of the switch. For detailed instructions, see the user's guide for your switch.
- Ensure that the Java Runtime Environment (JRE), version 1.2.2 or later, is installed on the server or workstation that will be used to manage the switch.
- Ensure that zoning is used to set up a SAN with heterogeneous hosts. Because the worldwide name is subject to change when the host is restarted, specify members by port name when you set up a zone.

Chapter 6. Setting up the storage subsystem in a heterogeneous environment

Follow the steps in this chapter to set up and run a heterogeneous storage subsystem that is managed by Storage Manager 8.21.

1. Ensure that your network has been configured and Storage Manager 8.21 is already running on one of the network-attached hosts as described in Chapter 2, "Installing storage management station software on AIX systems" on page 11, Chapter 3, "Installing storage management station software on HP-UX systems" on page 29, or Chapter 4, "Installing storage management station software on Solaris systems" on page 43.

   For switch zoning information, see "Setting IBM SAN switches" on page 63, or see your McData or INRANGE switch documentation.

2. Ensure that the following devices are powered up and ready:
   - All hosts
   - All fibre-channel switches
   - The storage management station
   - All FAStT controllers
   - All attached drive enclosures

3. If an AIX host is part of the network configuration that uses the Fast controller, then AVT/ADT, also known as the Automatic Logical Drive Transfer (ADT), must be disabled. See"Disabling AVT/ADT" on page 56.

   Attention: Because AVT/ADT must be disabled when running with an AX host and enabled when running with an HP-UX host, an AIX host and an HP-UX host cannot both be connected to the same FAStT server in a heterogeneous environment.

4. See the chapter about storage subsystem configuration in the *IBM FAStT Storage Manager Concepts Guide*.

   In the sections about the 3542 (for FAStT200), the 3552 (for FAStT500), or the 1742 (for FAStT700), find and perform these steps:
   a. Create logical drives
   b. Create storage partitions

   > **Note:** FAStT storage servers support up to 64 partitions and are shipped with storage partitioning enabled.

5. For each AIX host on the network, configure its partition by following these steps: Complete 9 on page 16to disable AVT/ADT. Complete 6 on page 15 to delete the access volume (Access Logical Drive)

   Attention: Only delete the Access Logical Drive from AIX hosts.

6. For each AIX host on the network, configure its partition by following the instructions starting at step 11 on page 16 in "Performing the initial configuration of storage subsystems on AIX hosts".

7. For each HP-UX host on the network, configure its partition by following the instructions starting at step 7 on page 35 in "Performing the initial configuration of storage subsystems on HP-UX hosts".

8. For each Solaris host on the network, configure its partition by following the instructions starting at step 7 on page 50 in "Performing the initial configuration of storage subsystems on Solaris hosts".

9. Each of the hosts should now be able to see their associated drive partition. Up to 64 hosts can be connected to each storage subsystem.

For more host-specific information, see the following chapters of this document:

- Chapter 2, "Installing storage management station software on AIX systems" on page 11
- Chapter 3, "Installing storage management station software on HP-UX systems" on page 29
- Chapter 4, "Installing storage management station software on Solaris systems" on page 43

Chapter 7. Storage Manager 8.21 with high-availability cluster services

The high-availability clustering services provided by Storage Manager 8.21 allow application services to continue when a hardware or software failure occurs. This system protects you from software failures as well as from the failure of a CPU, disk, or LAN component. If a component fails, its redundant partner component takes over cluster services and coordinates the transfer between components.

General information

Storage Manager 8.21 is certified for use with the following cluster services:
- MC/Service Guard versions A.11.09 and A.11.12 on HP-UX systems
- Veritas Cluster Server 1.3 and 2.0 on Solaris systems
- Veritas Volume Manager 3.1 and 3.2 on Solaris systems

This document does not describe how to install or configure cluster services. Refer to documentation provided with your cluster service products for this information.

Prerequisites for HP-UX

You can choose among many configurations when you set up clustering on an HP-UX system. A minimum configuration consists of two servers that are configured with both a primary and two standby LANs to establish a heartbeat LAN.

Provide fibre connections to the storage subsystem through two switches that provide the necessary redundant data path for the hosts. Ensure that each server has two HP Tachyon A5158A host bus adapters.

Prerequisites for Solaris and Veritas Cluster Server

The following sections contain general hardware requirements and additional information about the cluster services.

General hardware requirements

Each Solaris system in the cluster requires the following hardware:
- At least three Ethernet ports:
  - Two for the private network connections
  - At least one for the public network connection
- Two fibre host bus adapters for connection to the storage subsystem
- A SCSI connection for operating system disks
- Each Veritas Cluster Server system requires at least 128 MB of RAM and 35 MB of free disk space

System dependencies

This section provides information about Veritas Volume Manager, RDAC IDs, and single points of failure.

Veritas Volume Manager and Data Multipathing (DMP)

If you are using the Veritas Volume Manager, you must disable DMP either for all devices, or for just FAStT storage servers, depending on your version of the Veritas Volume Manager.

For Veritas Volume Manager 3.1 or earlier, follow these steps to disable DMP for all devices:

1. Open the /etc/system file in the vi text editor:

```
# vi /etc/system
```

2. Comment out the `forceload: drv/vxdmp` line.
3. Save and close the /etc/system file.

For Veritas Volume Manager 3.2 and later, you must disable DMP for FAStT storage servers only. Use option 17 and then option 5 of the submenu of the **vxdiskadm** utility to do so. Other devices can use DMP.

RDAC IDs

It is also recommended that you add up to eight additional IDs to the /etc/symsm/rmparams file. Complete the following steps to add them:

1. Open the /etc/symsm/rmparams file in the vi text editor:

```
# vi /etc/symsm/rmparams
```

2. Modify the `Rdac_HotAddIDs` line as follows:

```
Rdac_HotAddIDs:0:1:2:3:4:5:6:7:8
```

3. Save and close the /etc/symsm/rmparams file.

Single points of failure

When setting up cluster services, it is important to eliminate single points of failure because a single point of failure makes a cluster only as strong as its weakest component. Set up the storage subsystem for shared storage; for example, all the nodes in the cluster must recognize the same storage and the host types must be set correctly.

Appendix A. AIX system requirements

This appendix lists the minimum hardware and software requirements that an AIX system must meet to be used in a FAStT200, FAStT500, or FASTtT700 storage subsystem that is managed by Storage Manager 8.21.

Hardware requirements

The following hardware is required at a minimum.
- Any RS/6000 or eServer pSeries server that supports FC 6227 or FC 6228 host bus adapters
- 2 host bus adapters

Software requirements

The following software must be at the specified versions.
- AIX 4.3.3 or AIX 5.1
- Filesets associated with IBM RDAC driver installations listed in Table 11 or Table 12

RDAC installation requirements

The following filesets must be at the specified versions or later.
- For AIX 4.3.3:

Table 11. Filesets required for AIX 4.3.3 RDAC

PTF filesets	Version
devices.fcp.disk.array.diag	4.3.3.50
devices.fcp.disk.array.rte	4.3.3.82
devices.common.IBM.fc.rte	4.3.3.75
devices.pci.df1000f7.com	4.3.3.83
devices.pci.df1000f7.rte	4.3.3.75
devices.pci.df1000f9.rte	4.3.3.75
devices.scsi.scarray.rte	4.3.3.50

- For AIX 5.1:

Table 12. Filesets required for AIX 5.1 RDAC

PTF filesets	Version
devices.fcp.disk.array.diag	5.1.0.25
devices.fcp.disk.array.rte	5.1.0.29
devices.common.IBM.fc.rte	5.1.0.10
devices.pci.df1000f7.com	5.1.0.28
devices.pci.df1000f7.rte	5.1.0.15
devices.pci.df1000f9.rte	5.1.0.15
devices.scsi.scarray.rte	5.1.0.0

Note: The AIX RDAC driver files are not included on the installation CD; you must download them from the following Web site:

techsupport.services.ibm.com/server/fixes

For downloading instructions, see "Installing host software on AIX hosts" on page 13.

Appendix B. HP-UX system requirements

This appendix lists the minimum hardware and software requirements that an HP-UX system must meet to be used in a FAStT200, FAStT500, or FAStT700 storage subsystem that is managed by Storage Manager 8.21.

The HP-UX system must be an HP 9000/Series 800 server with:
- 180 MHz processor or faster
- 0.7 MB available on /opt and root (or root-equivalent) privileges for installing SMutil and SMagent
- 47.5 MB available disk space on /opt and at least 95 MB available on /tmp for installation in a Java Runtime Environment (JRE)

Ensure that the HP-UX host is running one of the following operating systems with the appropriate patches. Because patches can be superseded by more recent versions, refer to your operating system documentation or contact your operating system supplier to ensure that you have the correct patches.
- HP-UX 11.0 (32 bit or 64 bit) with the following patches:
  - B.11.00.53.2
  - B.11.00.53.8
  - PHKL_23939, version 1.0
  - PHKL_24004, version 1.0
  - PHKL_24027, version 1.0
  - PHKL_24612, version 1.0
  - PHKL_27003, version 1.0
  - To run HP-UX SDK for Java 2 Platform applications and applets using GUIs, you must also ensure that the HP C++ runtime libraries are installed on the HP-UX system; the latest version is available as patch PHSS_1658
- HP-UX 11.11 and HP-UX 11i with the following patches:
  - B.11.11.09
  - PHKL_25304
  - PHKL_25375
  - PHKL_25896
  - PHKL_PHKL_26519
- For high-availability clusters of HP 9000/Series 800 computers, install the HP MC/Service Guard software package

Ensure that the following maximum kernel parameters are configured, as shown in Table 13.

Table 13. HP-UX kernel parameter configuration requirements

Parameter	Description	Configuration
max_thread_proc 64	Maximum threads per process	1024
maxfiles	Soft-file limit per process	2048
maxuser	Influences other parameters	256 or greater
ncallout	Number of pending timeouts	4144

Appendix C. Solaris system requirements

This appendix lists the minimum hardware and software requirements that a Solaris system must meet to be used in a FAStT200, FAStT500, or FAStT700 storage subsystem that is managed with Storage Manager 8.21.

The Solaris system must be an Sparc S20 processor with:
- 256 MB system memory
- CD-ROM drive
- Mouse or similar pointing device
- Ethernet network interface card
- 1 MB available on /opt and root (or root-equivalent) privileges for installing RDAC

Ensure that the Solaris host is running one of the following operating systems with the appropriate patches. Because patches can be superseded by more recent versions, refer to your operating system documentation or contact your operating system supplier to ensure that you have the correct patches.
- Solaris 2.6 with the following patches:
  - 105181-32
  - 105356-19
  - 106429-02
- Solaris 7 with the following patches:
  - 106541-21
  - 107544-03
  - 108376-38
- Solaris 8 with the following patches:
  - 108528-15
  - 111111-03
  - 111310-01
  - 111293-04
- For high-availability clusters of Sparc S20 systems, install the Veritas Cluster Server software package

Appendix D. MC Service Guard configuration details

Note: The Hewlett Packard publication *Managing MC/ServiceGuard* suggests using **vgimport -m -s** with LVM commands during the configuration distribution process. It is important to consider, however, that the **vgimport** command does not preserve the primary and alternate paths of the imported volume groups.

When **vgimport** reconstructs the newly imported volume groups, entries are made in the /etc/lvmtab file. In this process of reconstruction, the system reorders disks in the file. The revised order of LUNs on the list causes LUNs to remain on non-preferred paths after failover. You might expect this condition to occur during distribution of the volume groups using the following LVM command:

```
vgimport -s -m /tmp/<vg_group_name>.map /dev/<vg_group_name>
```

This precaution applies only when you use the -s option to import devices or disks with redundant paths. The condition occurs because the -s option causes a search on the system for each disk when used with **vgimport**.

When the new list is prepared, the links that were previously designated as primary and alternate might not remain as they had been configured on node *A*.

Use the following procedure to correct the paths that were changed in this way when using **vgimport -m -s** with LVM commands.

Note: The following instructions are for a two node cluster. You can easily expand or adapt these instructions to accommodate a cluster which has more than two nodes.

1. Identify and note the primary and alternate paths for each device.

```
# SMdevices
```

   #SMdevices

2. Compare the primary and alternate paths for each device on node *A* to Node *B* by using the **vgdisplay -v** command. As previously noted, an expected behavior of the **vgimport -m -s** command is the changing of primary and alternate paths during export.

   On node *A*: **#vgdisplay -v <volume_group_name>**

   On node *B*: **#vgdisplay -v <volume_group_name>**

3. If the original primary path of a disk becomes an alternate path after the newly imported volume group entry is created in /etc/lvmtab, the order can be easily reverted by using **vgreduce** to remove the wrong primary path and then using **vgextend** to add the path back again:

```
#vgreduce
vg1 /dev/dsk/<device_name>#vgextend
 vg1 /dev/dsk/<device_name
```

Notices

This publication was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS ″AS IS″ WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publications. IBM may make improvements or changes (or both) in the product(s) or program(s) (or both), described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM
AIX
e(logo)
Netfinity
TotalStorage

Microsoft and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

IBM agreement for licensed internal code

You accept the terms of this Agreement[1] by your initial use of a machine that contains IBM Licensed Internal Code (called "Code"). These terms apply to Code used by certain Machines IBM or your reseller specifies (called "Specific Machines"). International Business Machines Corporation or one of its subsidiaries ("IBM") owns copyrights in Code or has the right to license Code. IBM or a third party owns all copies of Code, including all copies made from them.

If you are the rightful possessor of a Specific Machine, IBM grants you a license to use the Code (or any replacement IBM provides) on, or in conjunction with, only the Specific Machine for which the Code is provided. IBM licenses the Code to only one rightful possessor at a time.

Under each license, IBM authorizes you to do only the following:

1. execute the Code to enable the Specific Machine to function according to its Official Published Specifications (called "Specifications");
2. make a backup or archival copy of the Code (unless IBM makes one available for your use), provided you reproduce the copyright notice and any other legend of ownership on the copy. You may use the copy only to replace the original, when necessary; and
3. execute and display the Code as necessary to maintain the Specific Machine.

You agree to acquire any replacement for, or additional copy of, Code directly from IBM in accordance with IBM's standard policies and practices. You also agree to use that Code under these terms.

You may transfer possession of the Code to another party only with the transfer of the Specific Machine. If you do so, you must 1) destroy all your copies of the Code that were not provided by IBM, 2) either give the other party all your IBM-provided copies of the Code or destroy them, and 3) notify the other party of these terms. IBM licenses the other party when it accepts these terms. These terms apply to all Code you acquire from any source.

Your license terminates when you no longer rightfully possess the Specific Machine.

Actions you must not take

You agree to use the Code only as authorized above. You must not do, for example, any of the following:

---

1. Form Z125-4144

1. Otherwise copy, display, transfer, adapt, modify, or distribute the Code (electronically or otherwise), except as IBM may authorize in the Specific Machine's Specifications or in writing to you;
2. Reverse assemble, reverse compile, or otherwise translate the Code unless expressly permitted by applicable law without the possibility of contractual waiver;
3. Sublicense or assign the license for the Code; or
4. Lease the Code or any copy of it.

Glossary

This glossary includes terms for the IBM FAStT Storage Manager Version 8.21 for UNIX and AIX environments.

This glossary also includes terms and definitions from: *IBM Dictionary of Computing*. New York: McGraw-Hill, 1994.

The following cross-reference convention is used in this glossary:

See Refers you to (a) a related term, (b) a term that is the expanded form of an abbreviation or an acronym, or (c) a synonym or more preferred term.

Abstract Windowing Toolkit (AWT).. A Java graphical user interface (GUI).

access volume. A special logical drive that allows the host-agent to communicate with the controllers in the storage subsystem.

agent. A server program that receives virtual connections from the network manager (the client program) in an SNMP-TCP/IP network managing environment. The interface to a managed device.

auto volume transfer (AVT)/auto disk transfer (ADT). Provides automatic failover in case of controller failure on a storage subsystem.

AVT/ADT. See *auto volume transfer (AVT)/auto disk transfer (ADT)*.

AWT. See *Abstract Windowing Toolkit*.

bootstrap protocol (BOOTP). A TCP/IP protocol used by a diskless workstation or network computer to obtain its IP address and other network information such as server address and default gateway. Upon startup, the client station sends out a BOOTP request in a UDP packet to the BOOTP server, which returns the required information. Unlike RARP, which uses only the layer 2 (Ethernet) frame for transport, the BOOTP request and response use an IP broadcast function that can send messages before a specific IP address is known.

BOOTP. See *bootstrap protocol*.

client. A computer system or process that requests a service of another computer system or process that is typically referred to as a server. Multiple clients can share access to a common server.

dac. Disk array controller devices that represent the two controllers of the array.

dar. Disk array router that represents the entire array, including current and deferred paths to all LUNs (hdisks on AIX).

field replaceable unit. An assembly that is replaced in its entirety when any one of its components fails. In some cases, a field replaceable unit might contain other field replaceable units.

FRU. See *field replaceable unit*.

GUI. An acronym for graphical user interface.

hdisk. An AIX term; an hdisk represents a LUN on the array.

host. A system that is directly attached to the storage subsystem through a fibre-channel I/O path. This system is used to serve data (typically in the form of files) from the storage subsystem. (A system can be both a storage management station and a host at the same time.)

host computer. See *host*.

internet protocol (IP). In the Internet suite of protocols, a connection-less protocol that routes data through a network or interconnected networks and acts as an intermediary between the higher protocol layers and the physical network.

IP. See *internet protocol*.

JRE. An acronym for Java runtime environment.

LAN. See *local area network*.

logical unit number (LUN). An address that is used by the host to access a particular logical drive.

local area network (LAN). A computer network located on a user's premises within a limited geographic area.

LUN. See *logical unit number*.

man pages. Manual pages. UNIX command. To get information about UNIX commands, type `man <command_name>`; for example, to get information about the **ls** command, type `man ls`.

management information base (MIB). The information that is on an agent. It is an abstraction of configuration and status information.

MIB. See *management information base*.

NMS. An acronym for network management station.

NVSRAM. An acronym for nonvolatile storage random access memory.

ODM. An acronym for Object Data Manager.

PDF. An acronym for portable document format.

program temporary fix (PTF). A temporary solution or bypass of a problem diagnosed by IBM in a current unaltered release of a program.

PTF. See *program temporary fix*.

RDAC. See *redundant dual active controller*.

redundant dual active controller (RDAC). Also known as redundant disk array controller, provides a multipath driver for a storage subsystem. Used with AIX and Solaris hosts.

router. A computer algorithm that determines the path of least resistance for traffic on a computer network.

SAM. System Administrator Manager.

SAN. See *storage area network*.

storage area network (SAN). Dedicated, high-speed networks of directly connected storage elements designed to move large amounts of data between host-independent distributed storage devices.

storage management station. A system that is used to manage the storage subsystem. This system does not need to be attached to the storage subsystem through the fibre-channel I/O path.

server. A functional hardware and software unit that delivers shared resources to workstation client units on a computer network.

simple network management protocol (SNMP). A TCP/IP protocol that generally uses the user datagram protocol (UDP) to exchange messages between a management information base and a management client residing on a network. Because SNMP does not rely on the underlying communication protocols, it can be made available over other protocols, such as UDP/IP. In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application layer protocol. Information on devices that are managed is defined and stored in the application's management information base (MIB).

SNMP. See *simple network management protocol* and *SNMPv1*.

SNMPv1. The original standard for SNMP is now referred to as SNMPv1.

subnet. A part of a network that is identified by a portion of an IP address.

TCP. A communication protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol.

TCP provides a reliable host-to-host protocol between hosts in packed-switched communication networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

TCP/IP. See *transmission control protocol/internet protocol*.

transmission control protocol/internet protocol (TCP/IP). A set of communication protocols that support peer-to-peer connectivity functions for both local and wide-area networks.

trap. A mechanism for SNMP agents to notify the SNMP management station of significant events.

worldwide name (WWN). A unique identifier for a switch on local and global networks.

WWN. See *worldwide name*.

Index

A

about this document   xi
access volumes   2, 7
actions you must not take   78
address
  for e-mail comments   xiii
  for reader comment form   xiv
  of IBM   xiv
  of the IBM director of licensing   77
agent software
  installing   32, 48
  package   5
agreement for licensed internal code   78
AIX
  AVT/ADT, disabling   2
  client software
    installing   12
    requirements   12
  configuration   14, 17
  direct-attached configuration   3, 28
  firmware
    requirements   11
    upgrading   15
    verifying   13
  hardware requirements   11, 69
  host software
    installing   13
    prerequisites   13
  identifying
    controller ID numbers   22
    device names and bus numbers   23
    logical drives by operating system device
      names   25
  LUNs, redistributing in case of failure   27
  NVSRAM
    requirements   11
    upgrading   15
  operating system requirements   8
  prerequisites   12
  RDAC driver
    configuring   17
    installing   13, 17
    requirements   3, 13, 69
  requirements
    client software   12
    hardware   11, 69
    host software   13
    operating system   8
    RDAC driver   3, 13, 69
    software   69
    system   69
  restrictions   2, 3, 5, 7, 17
  software requirements   69
  system requirements   69
alert notifications, configuring   61
attributes
  dac   18, 21

attributes *(continued)*
  dar   18, 19
  definitions   19
  hdisk   18, 19, 21
  LUN   18, 19, 21
  RDAC driver   17
audience of this document   xi
AVT/ADT
  disabling   56
  enabling
    no multipath driver   56
    procedure   57
AVT/ADT, description   2

B

books   xii, 1
bus numbers and device names, identifying   23, 39

C

changing the host type   57
client software
  AIX
    installing   12
    requirements   12
  HP-UX
    installing   32
    requirements   30
  package   3, 4
  Solaris
    installing   45
    requirements   45
    uninstalling   51
cluster services
  configurations   58
  description   3
  hardware requirements   67
  HP-UX requirements   67
  MC/Service Guard   67
  Solaris requirements   67, 68
  system dependencies   67
  Veritas Cluster Server   67
  Veritas Volume Manager   67, 68
command line interface   2
commands
  fget_config   24
  lsattr   25
  lsdev   24
comments about this document, how to send   xiii
configuration details, MC Service Guard   75
configuration replication   2
configuring
  AIX   17
  alert notifications   61
  direct-attached   3, 28
  RDAC driver   17, 68

83

W

Readers' Comments — We'd Like to Hear from You

IBM FAStT Storage Manager Version 8.21 for
UNIX and AIX Environments
Installation and Support Guide

Publication No. GC26-7490-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

_____   _____
Name   Address

_____
Company or Organization

_____
Phone No.

IBM ®
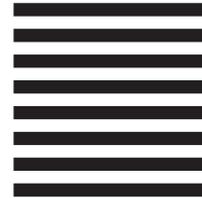
Printed in U.S.A.

Spine information:

IBM FAStT Storage Manager
Version 8.21 for
UNIX and AIX Environments

Installation and Support Guide

IBM